An Interop Publication

CONNEXIONS®

The Interoperability Report

May 1994 | *Special Issue: IP—The Next Generation* | Volume 8, No. 5

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, a division of ZD Expos, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

ISSN 0894–5926

From the Editor

The Internet faces one of its biggest challenges to date. Due to the incredible growth of the system in recent years, the 32-bit Internet address space is becoming depleted. In many ways this is analogous to "running out of phone numbers," a phenomenon quite common in many parts of the world. But unlike phone number format changes which can be very local in nature, the solution for the Internet will affect all connected hosts. Since routing and addressing is performed at the *Internet Protocol* (IP) layer, a new addressing structure means changes to IP itself. The problem has been recognized for quite some time, and several groups within the Internet Engineering Task Force (IETF) have been working on replacements for the current IP version 4 (IPv4). These efforts are collectively referred to as "IPng" or "IP: The Next Generation." In this edition we will look at several issues surrounding IPng, including the selection process, analysis of the problem at hand, and outlines of each of the IPng proposals.

Since IP is such a fundamental part of the Internet architecture, there can only be one IP—and only one IPng. Hence the IETF and the community at large must pick one of the existing IPng candidates (or perhaps develop yet another proposal). The IPng area has been formed within the IETF to consolidate all aspects of the selection process. The IPng area consists of a number of working groups plus a directorate. Our first article, by Phill Gross, IETF chairman, is an outline of the charge to the IPng area. Our second article, by Scott Bradner and Allison Mankin, describes the work of the directorate and the process leading up to the selection of IPng.

Each IPng proponent has been asked to provide a white paper which describes the technical details of their protocol and issues such as transition/coexistence, security, and support for new kinds of applications. The white paper guidelines are outlined in RFC 1550 which we include starting on page 11.

Exactly *when* IPng will be needed has been a matter of considerable debate. Obviously, we would like to have as much time as possible to develop and deploy an IPng which will last for many decades, but there is fear that the rapid address space depletion will force a quick (and premature) solution. Therefore, it is important to estimate how much time we really have. Frank Solensky gives an historical perspective on the current IPng efforts, followed by a presentation of what the address space depletion rates look like.

Next, each of the IPng proposals, CATNIP, TUBA and SIPP are described by the principal authors of each.

Finally, for the sake of completeness, we re-print an article which appeared in our September 1993 issue. The article is by Eric Fleischman and gives a user's perspective on IPng.

Charge to IPng area: A Direction for IPng

by Phill Gross, MCI Telecommunications

Introduction

At the Amsterdam IETF meeting, we held a BOF, entitled the "IP Decide BOF," on the process and progress of the IPng activities.

"IPng" stands for "IP: The Next Generation." The IPDecide BOF was chaired by Brian Carpenter. Minutes are available in the IETF directories, with the file name:

```
/ietf/93jul/ipdecide-minutes-93jul.txt.
```

Questions

The IPDecide BOF explored several facets of the IPng process, such as

- "What is the basis for choosing the next generation IP (i.e., what are the technical requirements and decision criteria)?"
- "With the advent of CIDR and new, more stringent address assignment policies, are we comfortable that we truly understand the level of urgency?"
- "Should the IETF or the marketplace make the final IPng decision?"

The BOF was held in a productive atmosphere, but did not achieve what could be called a clear consensus among the assembled attendees. In fact, despite its generally productive spirit, it did more to highlight the lack of a firm direction than to create it.

Consensus

The IPDecide BOF was followed the next evening by the open IESG plenary. During this session, the IESG and the assembled attendees discussed the IPng issues and seemed to arrive at a consensus based on the following set of bullets presented by the IETF chair:

- "The IETF needs to move toward closure on IPng." That is, the IETF should take active steps toward a technical decision, rather than waiting for the "marketplace" to decide.
- "The IESG has the responsibility for developing an IPng recommendation for the Internet community." That is, the IESG should provide leadership and take specific actions to help move the IETF toward a technical decision.
- "The procedures of the recommendation-making process should be open and published well in advance by the IESG."
- "As a part of the process, the IPng WGs may be given new milestones and other guidance to aid the IESG."
- "There should be ample opportunity for community comment prior to final IESG recommendation (e.g., there will be an extended Last Call)."

A direction for IPng

Building on this consensus, I'd like to announce a set of specific directions in the IESG that I hope will move us toward timely resolution of many of the key IPng issues.

The IESG will establish a temporary, *ad hoc,* "area" to deal specifically with IPng issues. The charter for this new IESG area is to develop a recommendation on which, if any, of the current proposals should be adopted as the "next IP." This recommendation will be submitted to the IESG and to the Internet community for review. Following an adequate period of review to surface any community concerns, the IESG will issue a final IPng recommendation. All of the current IPng-related working groups will be moved immediately into this new area.

This new area will be headed by two co-Area Directors (ADs) from within the IESG. I have asked Allison Mankin (NRL), current Transport Services AD, and Scott Bradner (Harvard), current Operational Requirements AD, to serve as co-ADs for this temporary area. I am very pleased to report that they have agreed to take this important assignment. (Because this is expected to be a temporary assignment, Scott and Allison will also continue to serve in their current IESG positions during this period.)

All IETF Areas are now expected to have Area Directorates. For the IPng Area, a Directorate will be especially important to bring additional viewpoints into the process. Therefore, I am asking that, as their first action, Scott and Allison form a specific IPng Directorate to act as a direction-setting and preliminary review body. The IPng process will continue to be completely open, and therefore reports and meeting notes from any IPng Directorate meetings will be published in a timely fashion.

Issues toward IPng resolution

Two important issues need resolution immediately before we can expect progress toward an IPng recommendation:

- *What is the scope of the effort?*

That is, should IPng be limited to solving the well-known scaling and address exhaustion issues; or should IPng also include advanced features such as resource reservation for real-time traffic?

The argument in favor of considering advanced features is that migration to a new IP is (hopefully, only!) a once-in-a-generation occurrence, and therefore all advanced features should at least be considered.

Arguments opposed to considering advanced features include the fact that we may not have time for this level of effort before the scaling and address exhaustion problems confront us, and that we may not have the necessary understanding and experience to make all the correct choices at this time.

- *What is the available timeframe?*

That is, before we can even begin to make an informed decision about the scope, we need a better understanding of the urgency and time constraints facing us.

Factors that affect the available time include the current rate of address assignments (which can give us an estimate of when we are currently projected to run out of addresses), the current policies governing address assignment (which can give us an understanding of how policies affect the assignment and utilization rates), the impact of CIDR aggregation, the development time for IPng, and the time needed to field and migrate to the new IPng.

Action items

Therefore, I am asking the new ADs and the Directorate to start immediately the following specific activities to help guide their ultimate IPng recommendation:

1. Develop an understanding of the available timeframe, covering at least the following issues:

- Review Internet growth metrics, such as the current address assignment and utilization rates. Develop an understanding of how the new address assignment policies impact the assignment and utilization rates.

continued on next page

A Direction for IPng *(continued)*

- Review the expected impact of CIDR address aggregation. Develop an understanding of the expected savings due to CIDR aggregation.
- Develop new technical guidelines for classless Internet addressing. Specific examples include guidelines for how to utilize variable length subnet masks, and how to utilize currently unused Class A & B addresses in a classless fashion in hosts and routers.
- Develop a strong understanding of the time required for the development, fielding, and migration for a new IP.
- Based on all the above issues,

  (a) develop an estimate for how long we have to to develop and deploy an IPng. This could be a set of estimates based on best/worst case estimates for how each of the above factors will affect the available timeframe.

  (b) Consider whether more stringent assignment policies might provide additional time. If so, recommend such policies.

  (c) make a recommendation on whether it is worthwhile to mount a serious effort to reclaim addresses and/or to renumber significant portions of the Internet.

2. Based on an informed judgment of the time constraints above, make a recommendation regarding the scope for IPng, i.e., should IPng consider scaling issues only or advanced topics also.

3. Based on the scope and time constraints, develop a clear and concise set of technical requirements and decision criteria for IPng. These should include, but not be limited to, the criteria outlined in the IESG statement (RFC 1380).

4. Based on the decision criteria, scope, and time constraints, make a recommendation on which of the current IPng candidates to accept, if any.

Report Finally, I am asking Scott and Allison to make a detailed report at the opening plenary of the next IETF meeting in November on the status of setting up their new area, and on their progress toward organizing the above work items. In particular, the status of the work items on timeframe should be fully reported. This will be followed by regular progress reports to the Internet community, at IETF meetings and in other appropriate forums.

Please join me in giving Scott and Allison our full cooperation, and in thanking them for accepting this daunting assignment. I feel confident that we will now make significant progress on the important IPng issues facing the Internet community.

—Phill Gross ,
IETF/IESG Chair

PHILL GROSS has recently joined MCI's new Data Services Division as the Director of Broadband Engineering. Among other activities, he is responsible for MCI's ATM service. Phill is a co-founder of the IETF and served as its chair for seven years until April 1994. He created the Internet Engineering Steering Group (IESG), that is now responsible for Internet standards decisions. In addition, he initiated such IETF institutions as Internet Drafts, working groups, Proceedings for each meeting, and the online IETF directories. He also co-chaired the Routing and Addressing (ROAD) group with Peter Ford (LANL). The ROAD group developed Classless Inter-Domain Routing (CIDR) and kicked off the BGP4 development. Prior to MCI, Phill was VP of network technology at ANS. Phill has an MS in computer science from Pennsylvania State University. In consideration of his long years of IETF service, his subscription to *MAD* magazine is kept up-to-date. He can be reached via Internet e-mail: `phill_gross@mcimail.com`

So, what's the problem?

The Internet suite of protocols ("TCP/IP") uses a 32-bit addressing scheme for all communication at the internet layer. This 32-bit address space is divided into several address *classes* as illustrated in Figure 1. Every Internet address can be thought of as a pair consisting of a network identifier (*netid*) and a host identifier (*hostid*). Since the 32-bit address length is fixed, it follows that the division between host- and netid can be made to "favor" either a large number of hosts on a given network (e.g., Class A) or a large number of networks with relatively few hosts on each (e.g., Class C). If we draw the class allocation as a pie chart (Figure 2), we see the consequences of the current address class structure. Notice, for instance, how Class A addresses take up 50% of the chart, but there are only 128 possible Class A networks, each with 16,777,215 hosts. The chart also gives the percentages of network addresses that have been "given out" by the *Internet Assigned Numbers Authority* (IANA). The use of *Classless Inter-Domain Routing* (CIDR) will serve as a short term solution by using the existing address space more efficiently, but IPng must be able to handle addresses larger than 32 bits. All of the IPng contenders were designed with this in mind. Additionally, new features such as security, support for mobile hosts, and real-time resource reservations are part of the requirements for IPng.

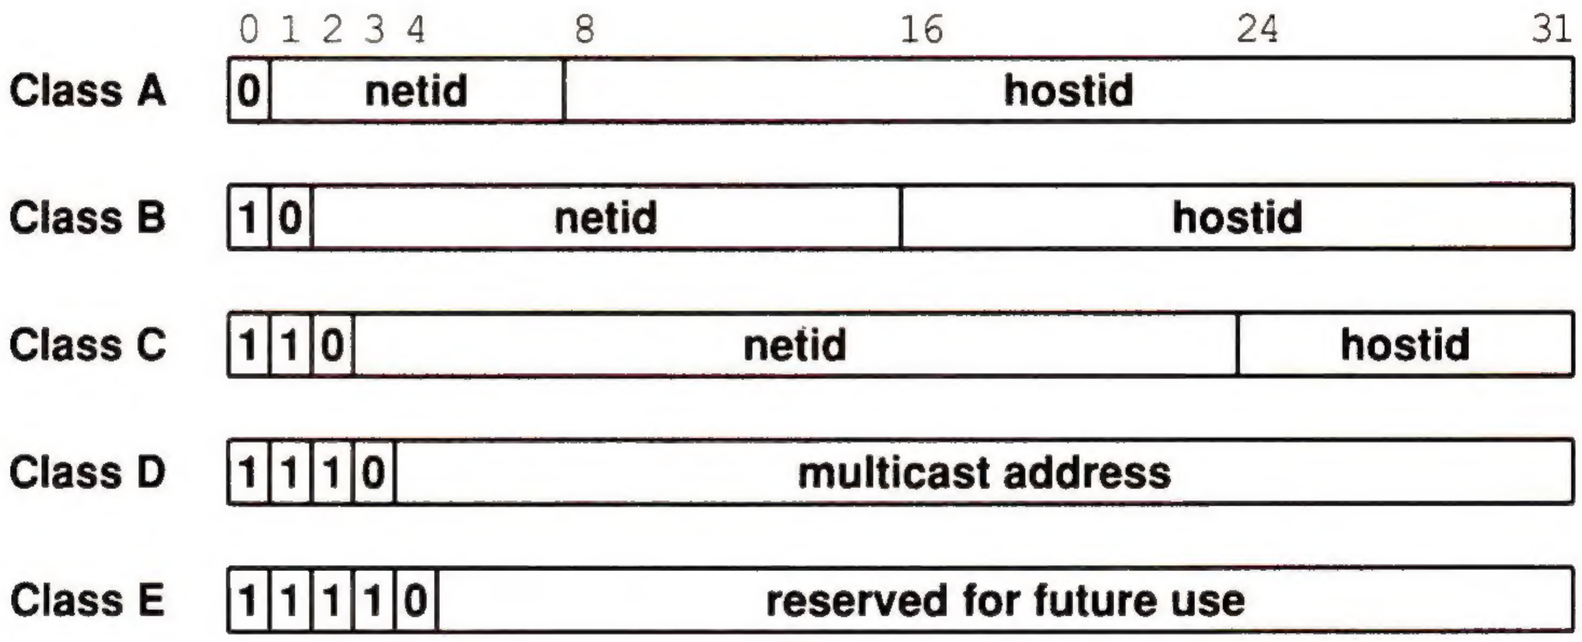

Figure 1 (Courtesy of Doug Comer): The five forms of IP addresses

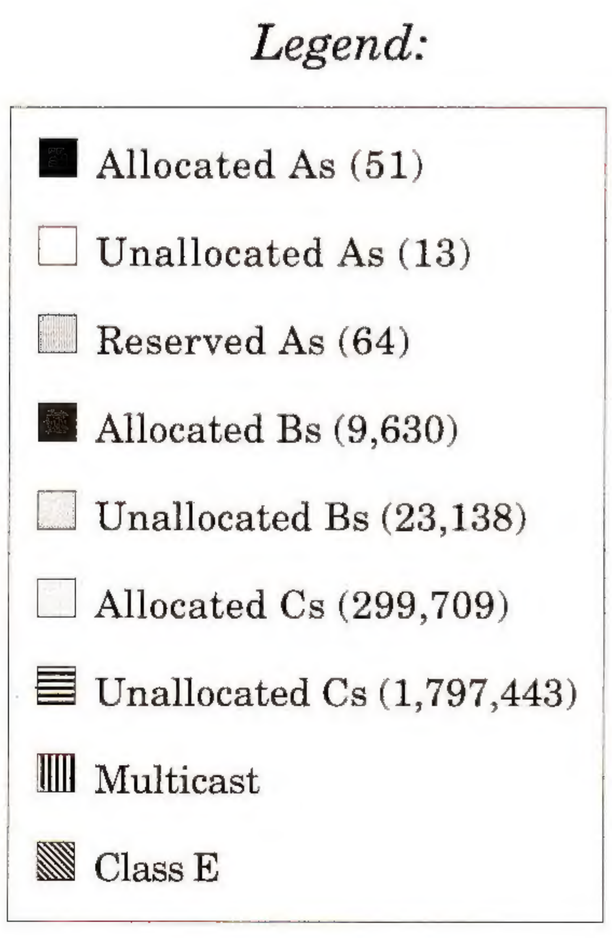

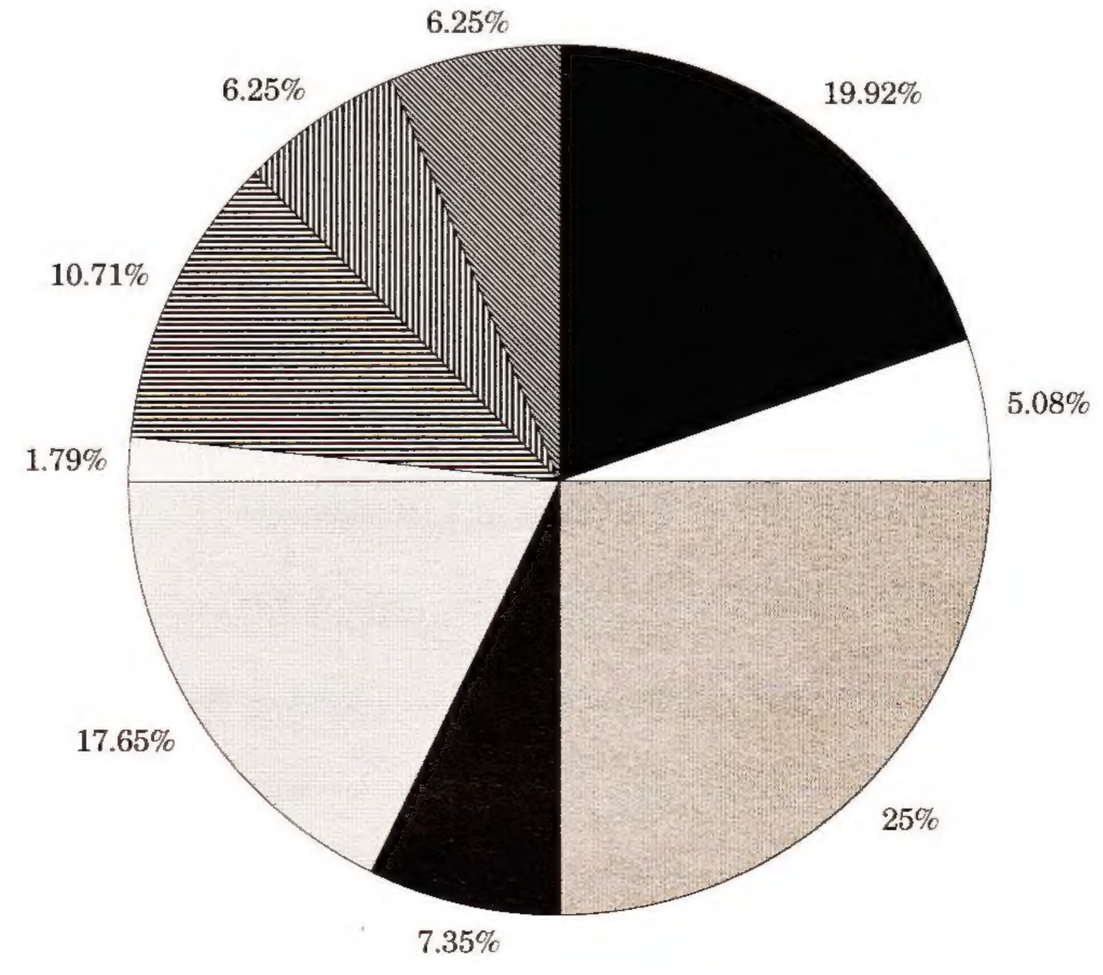

Figure 2 (Courtesy of Frank Solensky): The current picture

The IPng Selection Process

by
Scott Bradner, Harvard University
and
Allison Mankin, Naval Research Laboratory

Introduction

A new area was formed by the Internet Engineering Steering Group (IESG) on 7 September 1993 to consolidate all aspects of the process of selecting a replacement for IPv4. This new area was designated as the *IP: The Next Generation,* or *IPng,* area. It is a temporary area which will be dissolved after the selection process has been completed. We were asked to assume the duty of Area Directors on this new area (while not being excused from our existing duties as the Area Directors of the transport and operational requirements areas.) The working groups responsible for the development of the three current proposals for IPng were moved into the IPng area. They will disband, return to the Internet area, or become part of an IPng deployment area, when the IPng temporary area finishes its work.

Formation of the IPng area

In the document "A Direction for IPng" (see page 2) the IESG chair presented a brief description of the IPDecide BOF and other IPng activities during the Amsterdam IETF meeting, followed by a set of general direction statements, the announcement of the formation of the IPng area, and a specific charge to the new area. This article describes how the co-directors of the IPng area have decided to fulfill this charge and to "develop a recommendation on which, if any, of the current proposals should be adopted as the 'next IP.' "

As with all IETF areas, we were given considerable latitude to define the structure of the IPng area and the process by which the selection will be made. Some parts of the area structure and process were dictated by the charge, including the requirement of a directorate and the adoption of a policy of fully open discussions. Following traditional IETF practice, new working groups have been established to focus on specific aspects of the selection process. In addition we have established an expert panel, and are soliciting white papers to be used during the determination of the requirements for IPng and during the deliberation process.

White Papers

In a new process for the IETF, the IPng area is inviting the submission of *white papers* from the wider networking community (see RFC 1550, page 11). The papers fall into two categories: they can help define the requirements for an IPng or they can offer suggestions or solutions to these problems. The white papers are used as resource material for the working groups, as an information repository, and as a part of the permanent record of the IPng effort.

Some white papers have been specifically solicited from directorate members to indicate biases and opinions, from people active on `big-internet`, from IETF and non IETF researchers in the data networking field, from specific businesses, and from industry groups. The white paper process is open to all and papers have been received from a wide range of individuals. In addition, overviews of the specific IPng proposals were requested in the same form as the white papers.

Each of the white papers will be reviewed one or more times by members of the IPng directorate. The first review is to ensure that the ideas are presented completely and clearly. This review explicitly does not include a technical evaluation of the specific suggestions in the white paper. The results of the clarity review are then forwarded to the author(s).

If the author(s) wish, a revised version of the paper can be submitted. If this is done, the new version will replace the older one. At this point the papers are made available to the community as Internet Drafts. The Internet Drafts will then be reviewed for technical feasibility by members of the directorate and by the review panel. As is normal with Internet Drafts, the Internet community in general is expected to read and comment on the documents. After any revisions, the papers will be then reissued as Informational RFCs unless withdrawn by the author(s).

Insure the best proposals

A similar review process has been set up to ensure that the proposals for IPng are as good as they can be. An IPng choice should be based on a technical evaluation of the proposal and not be influenced by unclear or incomplete specifications. Using the same procedures established for the white papers, each of the proposal documents will be first reviewed for clarity and completeness with the reviewers giving specific suggestions for improvement. Once the documents have passed muster in this phase, they will be reviewed for technical feasibility. Note that this technical review is done within the context of the proposal, that is, reviewers cannot request changes just because of a disagreement over the width of the address, for example.

Open process

The entire IPng process is as open as we can make it. We do not want any hint that this important choice was made in secret by some unknown group. Minutes are kept during all of the directorate meetings and placed in the IPng public archives along with the archives of the directorate mailing list. In addition, the directorate will hold open meetings during the regular IETF meetings, and it will offer one or more MBONE meetings, as well (one took place January 25, 1994).

All of the IPng white papers and other documents are public documents except for the initial pre-clarity review version of the white papers. Using the normal IETF Internet Drafts process ensures the public view of the development of requirements via white papers.

Community input

Just as the IPng process can not be executed in private, it must not be done by a group isolated from the ideas and concerns of the data networking community. The white paper solicitation, the open review process, open directorate meetings, open working group meetings and mailing lists all are part of an ongoing effort to keep the community informed about the state of the selection process, the assumptions and priorities being used, and to give the community a basis on which to comment on the process.

IPng Working Groups

All of the existing working groups involved in the development of the IPng proposals were moved into the IPng area. During the development of the proposals some of the working groups have merged or changed names. The following illustration, courtesy of Stev Knowles, gives a genealogical history of the current working groups.

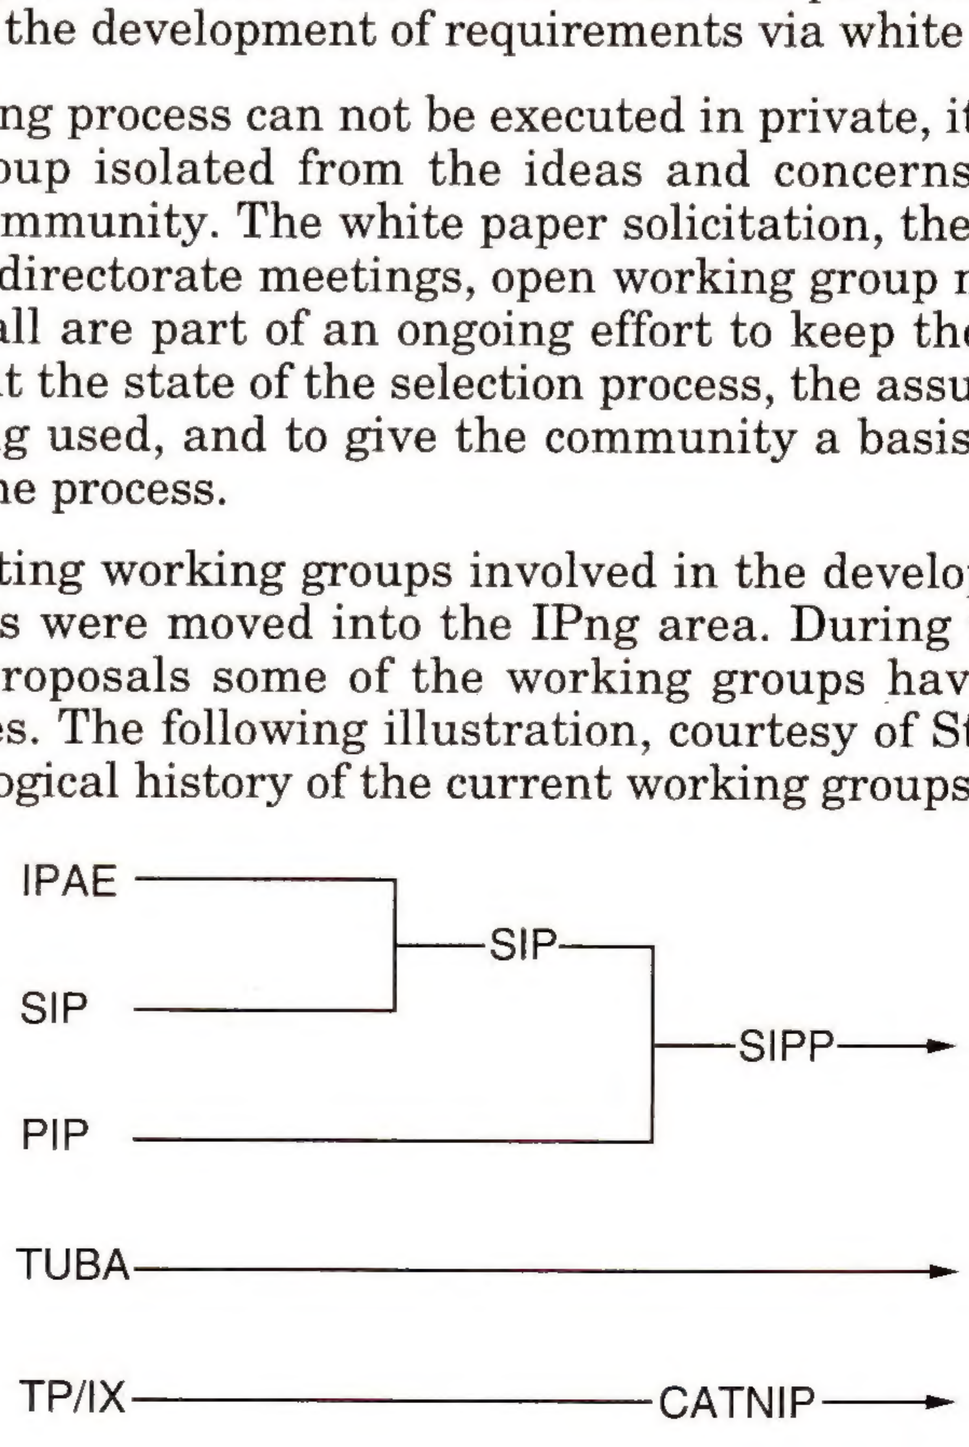

continued on next page

The IPng Selection Process *(continued)*

New IPng working groups

Three new working groups have been formed within the IPng area to develop information that will be used in the selection process and procedures that will be important after the selection is completed.

ALE

A new IETF working group, *Address Lifetime Expectations* (ALE), has been formed to make estimates of the remaining useful lifetime of the address space used by the existing version of IP. The estimate will be made taking into account the use of CIDR, the changing address assignment policies, and the availability of additional procedural documentation showing how to make more efficient use of assigned space. Frank Solensky (FTP Software) and Tony Li (Cisco) chair ALE.

IPng Requirements

A second group, *IPng Requirements* (IPNGREQ), is in place to complete the determination of the set of features and functions that a new IP should support. Since some of the desired features will require additional research and development, realistic estimates will be made for the availability timeframe for each of the features. This will be a short-running group with its major effort at the March IETF meeting in Seattle. Co-chairs are Jon Crowcroft (University College London) and Frank Kastenholtz (FTP Software).

Transition, coexistence and testing

A third working group, *Transition And Coexistence Including Testing* (TACIT), is in formation to develop an understanding of the operational issues involved in the migration of the Internet to a new internet protocol. There will be three chairs in all likelihood. Two have agreed to date, Atul Bansal (Digital) and Geoff Huston (AARnet).

Since the Internet must now be viewed as a utility and must continue to function during any transitions to new technologies, particular emphasis will be placed on planning a testing process. There may be bugs in the initial set of standards and almost certainly there will be interoperability problems with the initial implementations. It is very important that, by the time the new IP is deployed in a production network, it be as reliable as it can be. The IETF consensus is that IPng cannot be debugged in place. Since it is reasonable to expect that additional features will be added to IPng as time goes by (as features were added to the current IP), this group will be charged to create generic plans rather than target the existing proposals.

IPng Directorate

The charge to the IPng area included a requirement that the area have a directorate to act as a direction-setting and preliminary review body. We took some time to recruit this panel. We wanted to be sure that the people we selected were recognized technical aces but also that in addition to being able to articulate needs of corporation, customers, and community, they must be able to represent themselves. We did not want corporate mouth pieces.

We asked advice from many sources inside and outside of the IETF community while compiling the directorate. We wanted to insure that we had expertise in may areas including security, routing, international, national and regional network operations, large corporate networks, theoretical research, and protocol architecture, while insuring a depth of understanding on current IPng proposals.

The directorate that we recruited includes J. Allard (Microsoft), Steve Bellovin (AT&T), Jim Bound (Digital), Ross Callon (Wellfleet), Brian Carpenter (CERN), Dave Clark (MIT), John Curran (NEARnet), Steve Deering (Xerox), Dino Farinacci (Cisco), Paul Francis (NTT), Eric Fleischman (Boeing), Daniel Karrenberg (RARE), Mark Knopper (Ameritech), Greg Minshall (Novell), Paul Mockapetris (ISI), Rob Ullmann (Lotus) and Lixia Zhang (Xerox).

The IPng directorate holds teleconferences about every two weeks. Minutes are kept for these teleconferences and the open directorate meetings that will be held during IETF meetings. The minutes are placed in the IPng public archives. We also ask that the IPng directorate members monitor and respond to the `big-internet` mailing list. We chose to use the `big-internet` list instead of creating a new list for the IPng effort since `big-internet` was well established and focused on the very issues that an IPng list would.

External review panel

The directorate provides the major review and quality control of the IPng Area results. At Dave Clark's suggestion, we added the external review panel as a way to get review and advice from a larger community. The selection of IPng occurs during a time of remarkable growth for the Internet, in terms of its using population, markets, and potential applications. The expert review panel will offer views and input from individuals belonging to a wide range of industries, the cable industry, the classic telecommunications industry, and so on. It will also be the source of reviews by a key group of Internet thinkers, who for one reason or another did not join the directorate, but whose perspectives we did not want to risk not having. The responsibilities of the panel are to review the requirements document and the recommendation document.

Note that the area co-directors view the entire IETF as the internal review panel. We have solicited input and review from many folks to date. We welcome the entire community to pitch in!

IPng process

The ALE and IPng requirements working groups are scheduled to produce final reports after the Seattle IETF meeting (in late March 1994). At that time these reports will be combined to determine the final selection criteria. The proponents of each of the proposals will be asked to produce a white paper detailing how their proposal will meet the requirements and the associated timeframes. The proposal white papers will be reviewed by the IPng directorate and review panel and public comment will be invited.

A draft of the final IPng selection, with whatever specific suggestions may be warranted, will be produced by the area co-directors. This draft will be reviewed by the IPng directorate and the review panel. The area co-directors will take the results of this review into account to produce a final recommendation. This recommendation is currently due to be presented at the IETF meeting in July 1994.

The recommendation will be forwarded to the IESG for their consideration after an extended public review period. The IESG has stated that it will, if there is a suitable outcome, advance the selected protocol to Proposed Standard, and leave the other proposals as Experimental Standards.

IPng selection criteria

We are determined that politics not play a significant part in the IPng decision process. This may not be an easy separation to make for some people, but we want the replacement for IP to be selected by performing a technical evaluation of the proposals in relation to the technical requirements generated during the IPng process. We may have to explicitly deal with some political issues, such as ownership of base documents at some point in the process, but any issues of this type will be kept subservient to the technical evaluation.

continued on next page

The IPng Selection Process *(continued)*

Summary

We have established a procedure that we believe will result in the development of a common understanding of what the requirements for an IPng should be. This will produce a foundation upon which we can perform a careful analysis of the best way to meet these requirements. In addition to making the correct selection we must be sure that there is a solid understanding of how to test and deploy this technology. Now that the Internet is a utility, it must continue to function and function well even during the transition to the technology that will support our needs far into the future.

The IPng process cannot be one where the future is decided by some small bunch of people, no matter how right-thinking and well-intentioned. This process will only work if the community as a whole takes part. All of you are urged to read and comment on the white papers, the directorate mailing list and the requirements documents. All of you are urged to participate!

[This is an expanded version of a column that was published in *Network World* on November 22nd 1993, page 14.]

References

[1] V. Fuller, T. Li, J. Yu, K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[2] S. Deering, "Simple Internet Protocol Plus (SIPP) Specification," Internet Draft, work in progress.

[3] R. W. Callon, "TCP and UDP with Bigger Addresses (TUBA), A simple proposal for Internet addressing and routing," RFC 1347, June 1992.

[4] R. Ullmann, "TP/IX: The Next Internet," RFC 1475, June 1993.

[5] B. Carpenter, "IPng White Paper on transition and other considerations," Internet Draft, work in progress, March 1994.

[6] F. Kastenholz, C. Partridge, "Technical Criteria for Choosing IP: The Next Generation (IPng)," Internet Draft, work in progress, March 1994.

[7] S. Bellovin "Security Concerns for IPng," Internet Draft, work in progress, March 1994.

At NetWorld+Interop 94 in Las Vegas, don't miss the IPng session (C17) on Wednesday, May 5th at 3:30pm. The session is chaired by Scott Bradner.

SCOTT BRADNER has been involved in the design, operation and use of data networks at Harvard University since the early days of the ARPANET. He was involved in the design of the Harvard High-Speed Data Network (HSDN) and the Longwood Medical Area network (LMAnet) and is a founder of the New England Academic and Research network (NEARnet). He is currently chair of the technical committees of the LMAnet and NEARnet. He is the Operational Requirements area director and co-director of the temporary IPng area for the IETF, and is a member of the Internet Engineering Steering Group (IESG). In addition he was elected in 1993 as a trustee of the Internet Society. Mr. Bradner is a consultant at the Harvard Office for Information Technology, Network Service Division where he works on the design and development of network-based applications and extensions to the Harvard data network. He is also the director of the Harvard Network Device Test Lab and runs an annual series of router and bridge performance tests. He is a frequent speaker at technical conferences, a columnist for *Network World* and an instructor for Interop Company. E-mail: `sob@harvard.edu`

ALLISON MANKIN is part of a group at Naval Research Laboratory in Washington, DC that is developing all the pieces needed for the ATDnet or Wabitway, a gigabit research and service testbed of six federal agencies. She participated in the DARTnet and BLANCA testbeds as well. She has persistent interests in congestion control and quality of service. In addition to being co-Director of the temporary IPng Area, she is Director of the IESG Transport Services Area. She earned her Computer Science master's at Northeastern University in Boston. She is on the Editorial Board of *IEEE Network*. E-mail: `mankin@cmf.nrl.navy.mil`

IP: The Next Generation (IPng) White Paper Solicitation (RFC 1550)

by Scott Bradner and Allison Mankin

Status of this memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Introduction

The *IP: The Next Generation (IPng)* area in the IETF is soliciting white papers on topics related to the IPng requirements and selection criteria.

All interested parties are invited to submit white papers detailing any specific requirements that they feel an IPng must fulfill or any factors that they feel might sway the IPng selection. An example of the former might be a submission by a representative of a utility company detailing the scaling and addressing features which would be required to service future inclusion of utility meters on the network. An example of the other case might be a paper outlining the potential effect on IPng of some sections of the future network connectivity being provided via wireless networks.

At this time, we are not accepting white papers that evaluate specific IPng proposals. This type of document will be accepted after the various proposal documents are deemed to be clear and complete.

All white papers will be reviewed in a process described below. As a result of these reviews, each white paper will receive the focused attention of the IPng directorate and the community. The white papers will be used as resource materials by the IPng Area working groups, the directorate, the external review board and the Area Directors, during the selection process.

The deadline for the submission of these white papers is February 1, 1994, though early submission is encouraged. Submit white papers, general or topic questions, and so on, to `ipng-wp@harvard.edu`.

Document review process

All submitted documents will first be reviewed for clarity by members of the IPng directorate and the external review board. This review may produce suggestions to the author on areas of the document where there may be some confusion as to the meaning. Authors are urged to consider any such suggestions as constructive and to re-examine their text in light of the suggestions.

A separate technical review will then be done of the white paper. This review will be conducted within the context of the document. That is, the review still will not make value judgments on the white papers, but will assess technical feasibility. This review may also produce suggestions to the author.

The document will be submitted as an Internet Draft after these reviews have been completed and after whatever (if any) revisions the author decides to make. After a suitable period of time these documents will be submitted as informational RFCs unless withdrawn by the author. These documents will comprise a part of the historical record of the IPng process.

Document format requirements

All white papers must follow the format requirements listed in RFC 1543 and must not exceed 10 pages in length. They should not include the "status of memo" section; this will be added when the documents are posted as Internet Drafts. The reference version of the document must be in ASCII as is current practice with all RFCs.

continued on next page

IPng White Paper Solicitation *(continued)*

A *PostScript* version of the document may be submitted in addition to the ASCII version. (See RFC 1543 for the formatting procedures to use with *PostScript* documents.)

White Papers outline

This section details the white paper outline to be followed by someone who would like to express an opinion about the various factors involved in the IPng definition and selection process. Since these documents will be used as resource material by the various IPng working groups, the directorate, the external review board and the area directors, they should be well-focused and give specific references to data supporting their points.

Each white paper should begin with an executive summary of the important points of the document. This executive summary should not exceed 1/2 page in length.

The white paper should then address the issue or issues that the author feels should be understood during the IPng process. The total document should not exceed 10 pages in length. An author may submit more than one white paper if he or she feels that the level of detailed discussion on each topic warrants it.

Engineering considerations

In past discussions the following issues have been raised as relevant to the IPng selection process. This list is in no particular order. Any or all of these issues may be addressed as well as any other topic that the author feels is germane, but do not exceed the 10 page limit, please.

- *Scaling:* What is a reasonable estimate for the scale of the future data networking environment? The current common wisdom is that IPng should be able to deal with 10 to the 12th nodes.
- *Timescale:* What are reasonable time estimates for the IPng selection, development and deployment process or what should the timeframe requirements be? This topic is being evaluated by the ALE working group and a copy of all white papers that express opinions about these topics will be forwarded to that group.
- *Transition and deployment:* Transition from the current version to IPng will be a complex and difficult process. What are the issues that should be considered? The TACIT working group will be discussing these issues and a copy of all white papers that express opinions about these topics will be forwarded to that group.
- *Security:* What level and type of security will be required in the future network environment? What features should be in an IPng to facilitate security?
- *Configuration, administration and operation:* As networks get larger and more complex, the day-to-day operational aspects become ever more important. What should an IPng include or avoid in order to minimize the effect on the network operators?
- *Mobile hosts:* How important is the proliferation of mobile hosts to the IPng selection process? To what extent should features be included in an IPng to assist in dealing with mobile hosts?
- *Flows and resource reservation:* As the data networks begin to get used for an increasing number of time-critical processes, what are the requirements or concerns that affect how IPng should facilitate the use of resource reservations or flows?

- *Policy based routing:* How important is policy based routing? If it is important, what types of policies will be used? What requirements do routing policies and potential future global architectures of the Internet bring to IPng? How do policy requirements interact with scaling?
- *Topological flexibility:* What topology is anticipated for the Internet? Will the current general topology model continue? Is it acceptable (or even necessary) to place significant topological restrictions on the interconnectivity of networks?
- *Applicability:* What environment/marketplace do you see for the application of IPng? How much wider is it than the existing IP market?
- *Datagram service:* Existing IP service is "best effort" and based on hop-by-hop routed datagrams. What requirements for this paradigm influence the IPng selection?
- *Accounting:* How important a consideration should the ability to do accounting be in the selection of an IPng? What, if any, features should be included in an IPng to support accounting functions?
- *Support of communication media:* IPv4 can be supported over most known types of communications media. How important is this same flexibility to an IPng?
- *Robustness and fault tolerance:* To the extent that the Internet built from IPv4 has been highly fault tolerant, what are ways that IPng may avoid inadvertent decrease in the robustness (since some things may work despite flaws that we do not understand well)? Comment on any other ways in which this requirement may affect the IPng.
- *Technology pull:* Are there technologies that will pull the Internet in a way that should influence IPng? Can specific strategies be developed to encompass these?
- *Action items:* Suggested charges to the directorate, working groups or others to support the concerns or gather more information needed for a decision.

Security considerations

This RFC raises no security issues, but does invite comment on the security requirements of IPng.

Authors' addresses

Scott Bradner
Harvard University
10 Ware St.
Cambridge, MA 02138
Phone: +1 617 495-3864
E-mail: sob@harvard.edu

Allison Mankin
Naval Research Laboratory
c/o Code 5591
Washington, DC 20375–5000
Phone: +1 202 404-7030
E-mail: mankin@cmf.nrl.navy.mil

[Ed.: This RFC has been edited slightly to conform to the *ConneXions* "house style."]

CIDR Effects: Getting More out of IPv4

by Frank Solensky, FTP Software, Inc.

Introduction

Before we can decide upon the features and functionality that the eventual IPng protocol will provide, we need to estimate the time frame that is available to us during which the decision, design, deployment and transition towards this new protocol must all take place. Without having some idea of what the timeframe is, we run the risk of either deploying something too quickly and not sufficiently flexible to handle future needs or spending too much time on the design phase, making the transition and deployment a panic-driven event.

This article presents some of the history that has led to the current IPng efforts, followed by a heavily qualified presentation of what the growth rates look like.

The Class B numbering problem

Around 1990, some were becoming concerned that the assignment of IPv4 network numbers was becoming too concentrated within the Class B address space—those network numbers containing a 14-bit network number identifier and a 16-bit host number [3,12]. This was based on the realization that, not unlike Goldilock's assessment of the Bear family's beds, the size of the Class B network number space was best suited for most organizations. At that time,.approximately 20% of the available Class B network numbers were already allocated; if it were to continue doubling about every 14 months as it had up to that point, we would have exhausted the Class B address space within a few weeks of the time this issue went to press.

In order to avoid depleting the Class B address space completely, it was decided that the NIC would severely limit the assignment of these addresses to those organizations that could clearly demonstrate the need for it, otherwise the organization would instead be assigned a block of consecutive Class C network numbers [14].

The routing table problem

This action had the desired effect of deflecting some of the growth of the network number assignments into the Class C address space and simultaneously making more effective utilization of the total IPv4 address space, as illustrated in Figure 1. The trend lines in this diagram show the "best-fit" logistic curve through the data since the beginning of 1992. A brief description and application of the logistic curve in predicting future growth of networks can be found in [13] and [15].

While this took care of the immediate problems in the Class B address space, it was also recognized that this strategy would cause the routing tables to grow at a much faster rate than before [7]. This is due to the fact that storing multiple Class C network numbers required proportionally greater resources than a single Class B number. All of the nodes within an organization's network could no longer be referenced with a single entry in a routing table: each of the individual network numbers assigned to the organization had to be announced and stored separately.

CIDR

Recognition of the routing table explosion problem led to the development of *Classless Inter-Domain Routing,* or CIDR as it is more commonly known [2,4,6]. As its name implies, CIDR breaks away from the traditional concept of dividing the address space into classes and incorporates network masks into the routing protocols.

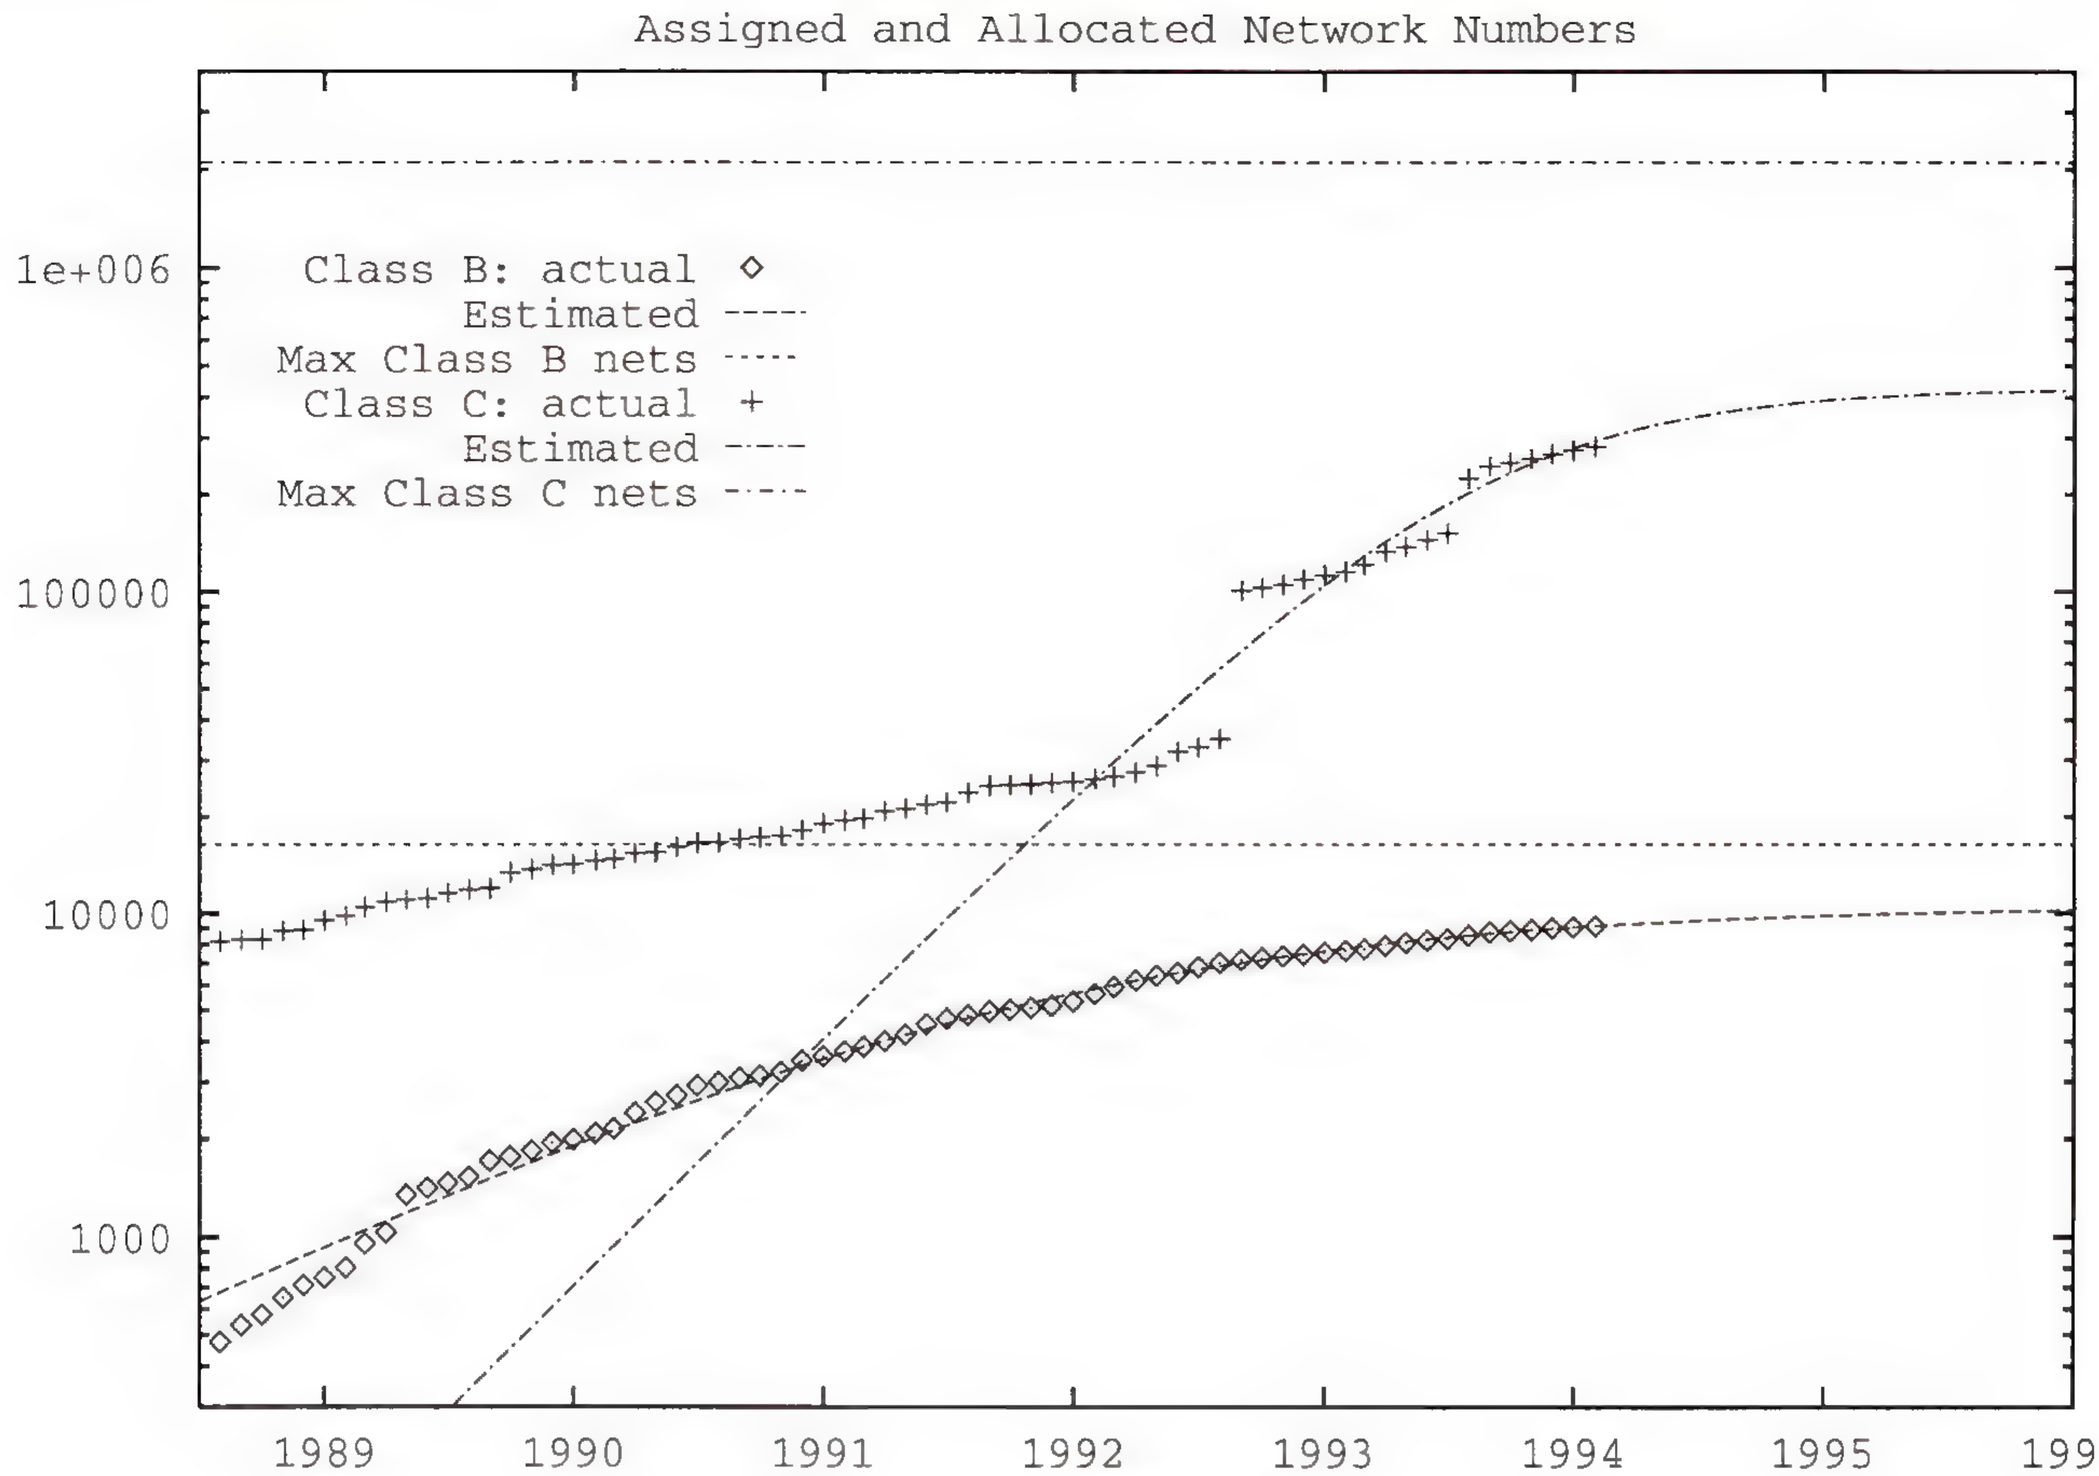

Figure 1

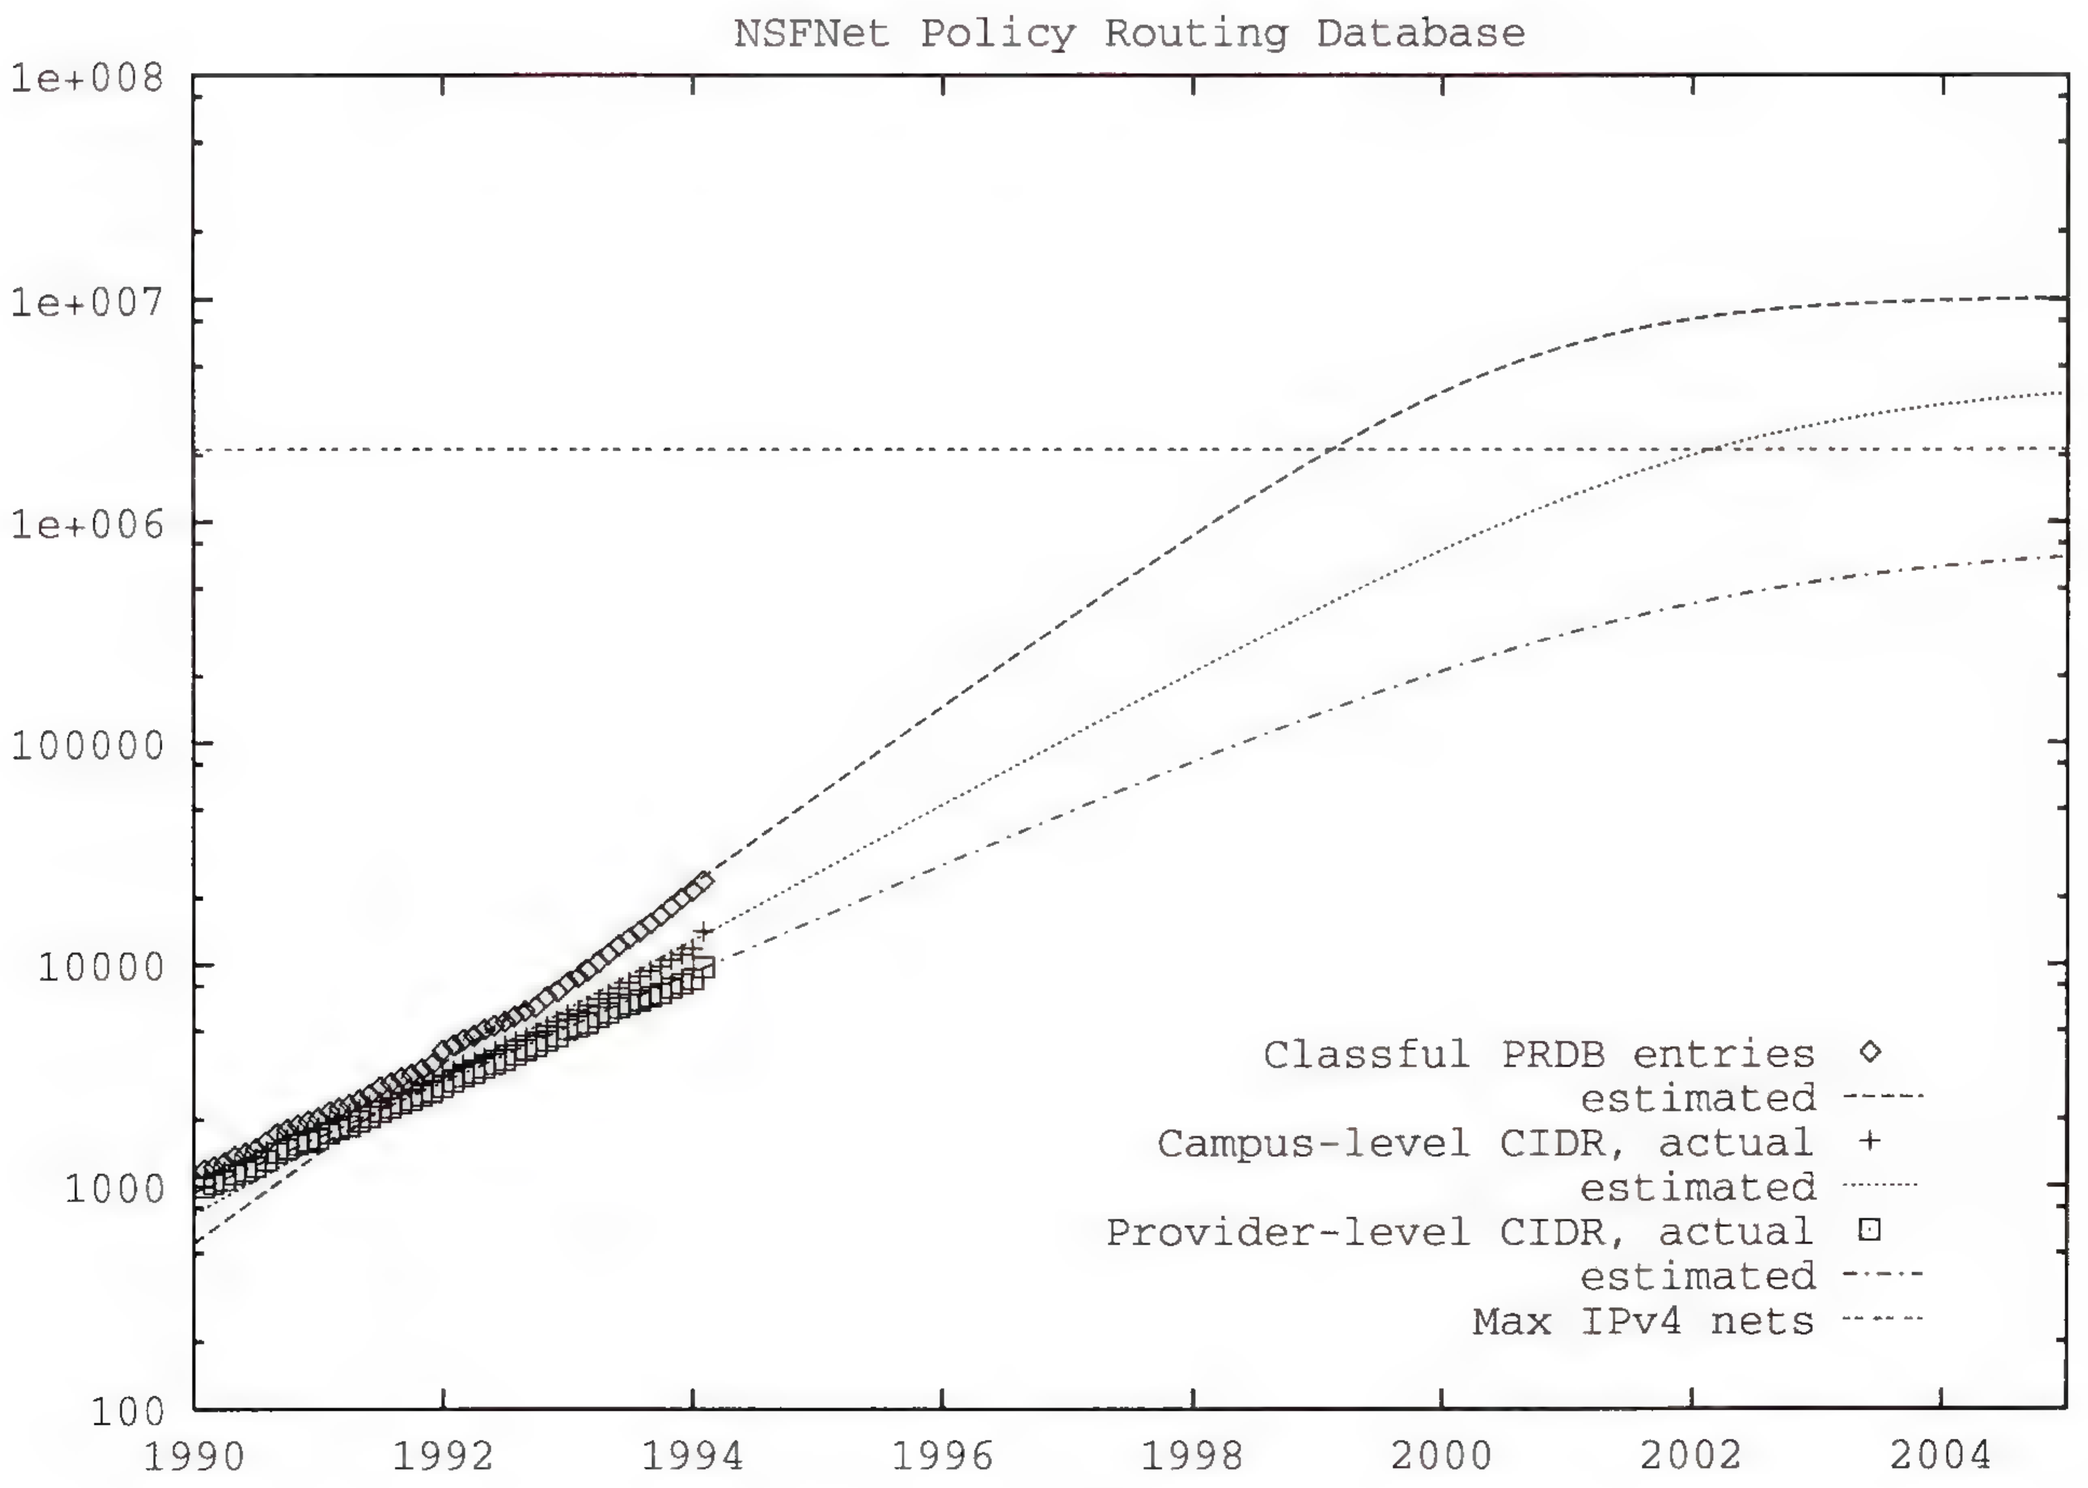

Figure 2

continued on next page

CIDR Effects *(continued)*

CIDR also goes one step further to make routing more efficient by aggregating consecutively numbered networks into a single announcement. Rather than just merging all of the network numbers assigned to an organization into a single routing table entry, this means that a single routing table entry could now announce multiple organizations simply by announcing a less-specific network mask.

OSPF [10], already capable of announcing individual subnetworks, needed only minor implemenation changes. Network masks were added to BGP version 4 [11] and RIP version 2 [9] so that they could also aggregate network numbers into a single routing update announcement.

Figure 2 provides a very rough approximation on how much of an effect CIDR would have had on the size of the routing tables had it already been deployed, then extrapolates the resulting timeseries off into the future. The top curve on this graph depicts the status quo: blocks of addresses continue to be announced as separate entries. By this reckoning, routers would need to keep track of one hundred thousand network number entries in mid- to late-1995. By deploying CIDR in a manner so that the aggregation occurs only at the network level, this gets pushed off towards the beginning of 1997. With full deployment of CIDR, current trends suggest that routers wouldn't need to keep track of one hundred thousand routes until the start of 1998. Another way to look at this would suggest that full CIDR deployment means that the size of the routing tables would not return to their current size until the start of 1996.

Your mileage may vary

Readers might then conclude that this evaluation proves that there's really nothing to worry about—CIDR saves IPv4 and the rest of this issue contains interesting technical exercises but little else of consequence. That conclusion would have to be reached by ignoring these other factors:

- The historic data which forms the basis of the trend line is determined by existing technologies. It does not attempt to anticipate either new technologies (such as wireless networking [5], ATM [8] or cable) or services (e.g., libraries [1] and utility companies) that may create additional demand for large slices of the remaining address space.
- These estimates do not reflect the potential impact of policy restrictions on the routing table size. This could limit the effectiveness of aggregating network numbers into wider masks if a network number in the middle of a block must traverse a different path.
- If the analysis is correct in suggesting that we may not be in imminent danger of the Internet collapsing under its increasing success, we can make use of the additional time to incorporate functionality into the basic structure that may be more difficult to add into the existing one.

The aim of the IPv4-ALE (*IPv4 Address Lifetime Expectation*) Working Group will be to develop a more concrete estimate for the time that remains before the IPv4 address space will no longer be able to satisfy the requirements of the growing Internet and to identify locations where the address space is being used inefficiently. Requests to join the working group's mailing list should be sent via e-mail to: **`ipv4-ale-request@ftp.com`**.

References

[1] Barron, B. "TRnet: A Possible Future Use of the Internet," *Connexions,* Volume 7, No. 12, December 1993.

[2] Braun, H-W., Ford, P., Rekhter, Y., "CIDR and the Evolution of the Internet Protocol," *ConneXions,* Volume 7, No. 9, September 1993.

[3] Chiappa, N., "The IP Addressing Issue," October 1990.

[4] Crocker, D., "The ROAD to a New IP," *ConneXions,* Volume 6, No. 11, November 1992.

[5] Dayem, R. A., "A Map to Wireless Networking and Mobile Computing," *ConneXions,* Volume 7, No. 9, September 1993.

[6] V. Fuller, T. Li, J. Yu, K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[7] Gross, P., Almquist, P., "IESG Deliberations on Routing and Addressing," RFC 1380, November 1992.

[8] Laubach, M., "ATM for your internet—But When?," *ConneXions,* Volume 7, No. 9, September 1993.

[9] Malkin, G., "RIP Version 2 Carrying Additional Information," Internet Draft, October 1993.

[10] Moy, J., "OSPF Version 2," Internet Draft, September 1993.

[11] ed. Rekhter, Y., Li, T., "A Border Gateway Protocol 4 (BGP-4)," Internet Draft, January 1994.

[12] Solensky, F., "Continued Internet Growth," Proceedings of the Eighteenth Internet Engineering Task Force, August 1990.

[13] Solensky, F., "The Growing Internet," *ConneXions,* Volume 6, No. 5, May 1992.

[14] Topolcic, C., "Schedule for IP Address Space Management Guidelines," RFC 1367, October 1992 (see also RFC 1467, August 1993).

[15] Gurbaxani, V. "Diffusion in Computer Networks: The Case of BITNET," *Communications of the ACM,* December 1990.

[16] Lottor, M., "Internet Growth (1981–1991)," RFC 1296, Jan. 1992.

[17] D. Clark, L. Chapin, V. Cerf, R. Braden, R. Hobby, "Towards the Future Internet Architecture," RFC 1287 December 1991.

[18] Z. Wang, J. Crowcroft, "Two-tier address structure for the Internet: A solution to the problem of address space exhaustion," RFC 1335, May 1992.

[19] B. Carpenter, "IPng White Paper on transition and other considerations," Internet Draft, work in progress, March 1994.

[20] F. Kastenholz, C. Partridge, "Technical Criteria for Choosing IP: The Next Generation (IPng)," Internet Draft, work in progress, March 1994.

FRANK SOLENSKY holds a BS and MS from New York University and has been working with internetworking software for most of the last nine years. He has been a member of the IETF since 1989 and recently joined FTP Software, working on Neat Stuff. He can be reached via the Internet at either `solensky@ftp.com` or at `bosox-request@world.std.com`

The CATNIP: Purrposed Common Architecture for the Internet

by Michael McGovern and Robert Ullmann

Introduction

This article describes a common architecture for the network layer protocol. The *Common Architecture for Next Generation Internet Protocol* (CATNIP) provides a compressed form of the existing network layer protocols. Each compression is defined so that the resulting network protocol data units are identical in format. The fixed part of the compressed format is 16 bytes in length, and may often be the only part transmitted on the subnetwork.

Robert Ullmann wrote the first description of a possible Internet Version 7 protocol in the summer and fall of 1989. Much of the thinking shown in this approach was paralleled by work done by Ross Callon under the name TUBA. The first version of TUBA was published in RFC 1347.

With some attention paid to details, it is possible for a transport layer protocol (such as TCP) to operate properly with one end system using one network layer (e.g., IP version 4) and the other using some other network protocol, such as CLNP. Using the CATNIP definitions, all the existing transport layer protocols used on connectionless network services will operate over any existing network layer protocol.

CATNIP uses *cache handles* to provide both rapid identification of the next hop in high performance routing as well as abbreviation of the network header by permitting the addresses to be omitted when a valid cache handle is available. The fixed part of the network layer header carries the cache handles.

The cache handles are either provided by feedback from the downstream router in response to offered traffic, or explicitly provided as part of the establishment of a circuit or flow through the network. When used for flows, the handle is the locally significant flow identifier.

When used for circuits, the handle is the Layer 3 peer-to-peer logical channel identifier, and permits a full implementation of network-layer connection-oriented service if the routers along the path provide sufficient features. At the same time, the packet format of the connectionless service is retained, and hop by hop fully addressed datagrams can be used at the same time. Any intermediate model between the connection-oriented and the connectionless service can thus be provided over cooperating routers.

CATNIP objectives

The first objective of CATNIP is a practical recognition of the existing state of internetworking, and an understanding that any approach must encompass the entire problem. While it is common in the IP Internet to dismiss CLNP with various amusing phrases, it is hardly realistic.

CATNIP integrates CLNP, IP, and IPX. The CATNIP design provides for any of the transport layer protocols in use, for example TP4, CLTP, TCP, UDP, IPX and SPX to run over any of the network layer protocol formats: CLNP, IP (version 4), IPX, and CATNIP.

Incremental infrastructure deployment

The best use of CATNIP is to begin to build a common Internet infrastructure. The routers and other components of the common system are able to use a single consistent addressing method, and common terms of reference for other aspects of the system.

CATNIP is designed to be incrementally deployable in the strong sense: you can plop a CATNIP system down in place of any existing network component and continue to operate normally with no reconfiguration. (Note: not "just a little." None at all. The number of "little changes" suggested by some proposals, and the utterly enormous amount of documentation, training, and administrative effort then required, astounds the present authors.) The vendors do all of the work.

There are also no external requirements; no "border routers," no requirement that administrators apply specific restrictions to their network designs, define special tables, or add things to the *Domain Name System* (DNS). When the end users and administrators fully understand the combined system, they will want to operate differently, but in no case will they be forced. Not even in small ways. Networks and end-user organizations operate under sufficient constraints on deployment of systems anyway; they do not need a new network architecture adding to the difficulty.

Typically deployment will occur as part of normal upgrade revisions of software, and due to the "swamping" of the existing base as the network grows. (When the Internet grows by a factor of 5, at least 80% will then be "new" systems.) The users of the network may then take advantage of the new capabilities. Some of the performance improvements will be automatic, others may require some administrative understanding to get to the best performance level.

The CATNIP definitions provide stateless translation of network datagrams to and from CATNIP and, by implication, directly between the other network layer protocols. A CATNIP-capable system implementing the full set of definitions can interoperate with any existing protocol. Various subsets of the full capability may be provided by some vendors.

No address translation

Note that there is no "address translation" in the CATNIP specification. (While it may seem odd to state a negative objective, this is worth saying as people seem to assume the opposite.) There are no "mapping tables," no magic ways of digging translations out of the DNS or X.500, no routers looking up translations or asking other systems for them.

Addresses are modified with a simple algorithmic mapping, a mapping that is no more than using specific prefixes for IP and IPX addresses. Not a large set of prefixes; one prefix. The entire existing IP version 4 network is mapped with one prefix and the IPX global network with one other prefix. (The IP mapping does provide for future assignment of other IANA/IPv4 domains that are disjoint from the existing one.) This means that there is no immediate effect on addresses embedded in higher level protocols.

Higher level protocols not using the full form (those native to IP and IPX) will eventually be extended to use the full addressing to extend their usability over all of the network layers.

No Legacy Systems

CATNIP leaves no systems behind: with no reconfiguration, any system presently capable of IP, CLNP, or IPX retains at least the connectivity it has now. With some administrative changes (such as assigning IPX domain addresses to some CLNP hosts for example) on other systems, unmodified systems may gain significant connectivity. IPX systems with registered network numbers may gain the most.

continued on next page

CATNIP *(continued)*

Limited scope

CATNIP defines a common network layer packet format and basic architecture. It intentionally does not specify ES–IS methods, routing, naming systems, auto-configuration and other subjects not part of the core Internet-wide architecture. The related problems and their (many) solutions are not within the scope of the specification of the basic common network layer.

Existing addresses and network numbers

The Internet's version 4 numbering system has proven to be very flexible, (mostly) expandable, and simple. In short: it works. However, there are two problems. Neither was considered serious when CATNIP was first developed in 1988 and 1989, but both are now of major concern:

- The division into network, and then subnet, is insufficient. Almost all sites need a network assignment large enough to subnet. At the top of the hierarchy, there is a need to assign administrative domains.
- As bit-packing is done to accomplish the desired network structure, the 32-bit limit causes more and more aggravation.

Another major addressing system used in open internetworking is the OSI method of specifying *Network Service Access Points* (NSAPs). The NSAP consists of an authority and format identifier, a number assigned to that authority, an address assigned by that authority, and a selector identifying the next layer (transport layer) protocol. This is actually a general multi-level hierarchy, often obscured by the details of specific profiles. (For example, CLNP doesn't specify 20 octet NSAPs, it allows any length. But various GOSIPs profile the NSAP as 20 octets, and IS–IS makes specific assumptions about the last 1–8 octets. And so on.)

The NSAP does not directly correspond to an IP address, as the selector in IP is separate from the address. The concept that does correspond is the NSAP less the selector, called the *Network Entity Title* or NET. (An unfortunate acronym, but one we will use to avoid repeating the full term.) The usual definition of NET is an NSAP with the selector set to 0; the NET used here omits the 0 selector.

There is also a network numbering system used by IPX, a product of Novell, Inc. (referred to from here on as Novell) and other vendors making compatible software. While IPX is not yet well connected into a global network, it has a larger installed base than either of the other network layers.

Network Layer address

The network layer address looks like:

The fields are named in the usual OSI terminology although that leads to an oversupply of acronyms. Here are more detailed descriptions of each field:

length: the number of bytes (octets) in the remainder of the address.

AFI: the Authority and Format Identifier. A single byte value, from a set of well-known values registered by ISO, that determines the semantics of the IDI field.

IDI: the Initial Domain Identifier, a number assigned by the authority named by the AFI, formatted according to the semantics implied by the AFI, that determines the authority for the remainder of the address.

DSP: Domain Specific Part, an address assigned by the authority identified by the value of the IDI.

Note that there are several levels of authority. ISO, for example, identifies (with the AFI) a set of numbering authorities (like X.121, the numbering plan for the PSPDN, or E.164, the numbering plan for the telephone system). Each authority numbers a set of organizations or individuals or other entities. (For example, E.164 assigns 16172477959 to me as a telephone subscriber.)

The entity then is the authority for the remainder of the address. I can do what I please with the addresses starting with (AFI=E.164) (IDI=16172477959). Note that this is a delegation of authority, and not an embedding of a data-link address (the telephone number) in a network layer address. The actual routing of the network layer address has nothing to do with the authority numbering.

The domain-specific part is variable length, and can be allocated in whatever way the authority identified by the AFI+IDI desires. (But note that things like GOSIPs and ES–IS as presently implemented put other, probably ill-advised, constraints on the DSP.)

Network Layer datagram

The common architecture format for network layer datagrams is described below. The design is a balance between use on high performance networks and routers, and a desire to minimize the number of bits in the fixed header. CATNIP avoids the mistake of making the fixed field too small. Using the current state of processor technology as a reference, the fixed header is all loaded into CPU registers on the first memory cycle, and it all fits within the operation bandwidth. The header leaves the remaining data aligned on the header size (128 bits); with 64 bit addresses present and no options it leaves the transport header 256 bit aligned.

On very slow and low performance networks, the fixed header is still fairly small, and could be further compressed by methods similar to those used with IP version 4 on links that consider every bit precious. In between, it fits nicely into ATM cells and radio packets, leaving sufficient space for the transport header and application data.

NLPID (70)	Header Size	D	S	R	M	E	MBZ	Time to Live
Forward Cache Identifier								
Datagram Length								
Transport Protocol				Checksum				
Destination Address...								
Source Address...								
Options								

NLPID: The first byte (the network layer protocol identifier in OSI) is an 8 bit constant 70 (hex). This corresponds to Internet Version 7.

Header Length: The header length is an 8-bit count of the number of 32-bit words in the header. This allows the header to be up to 1020 bytes in length.

continued on next page

CATNIP *(continued)*

Flags: This byte is a small set of flags determining the datagram header format and the processing semantics. The last three bits are reserved, and must be set to zero. (Note that the corresponding bits in CLNP version 1 are 001, since this byte is the version field. This may be useful.)

Destination Address Omitted: When the destination address omitted (DAO) flag is zero, the destination address is present as shown in the datagram format diagram. When a datagram is sent with an FCI that identifies the destination and the DAO flag is set, the address does not appear in the datagram.

Source Address Omitted: The source address omitted (SAO) flag is zero when the source address is present in the datagram. When datagram is sent with an FCI that identifies the source and the SAO flag is set, the source address is omitted from the datagram.

Report Fragmentation Done: When this bit (RFD) is set, an intermediate router that fragments the datagram (because it is larger than the next subnetwork MTU) should report the event with an ICMP "Datagram Too Big" message. (Unlike IP version 4, which uses DF for MTU discovery, the RFD flag allows the fragmented datagram to be delivered.)

Mandatory Router Option: The mandatory router option (MRO) flag indicates that routers forwarding the datagram must look at the network header options. If not set, an intermediate router should not look at the header options. (But it may anyway; this is a necessary consequence of transparent network layer translation, which may occur anywhere.) The destination host, or an intermediate router doing translation, must look at the header options regardless of the setting of the MRO flag. A router doing fragmentation will normally only use the RFD flag in options to determine whether options should be copied within the fragmentation code path. (It might also recognize and elide null options.) If the MRO flag is not set, the router may not act on an option even though it copies it properly during fragmentation. If there are no options present, MRO should always be zero, so that routers can follow the no-option profile path in their implementation. (Remember that the presence of options cannot be divined from the header length, since the addresses are variable length.)

Error Report Suppression: The ERS flag is set to suppress the sending of error reports by any system (whether host or router) receiving or forwarding the datagram. The system may log the error, increment network management counters, and take any similar action, but ICMP error messages or CNLP error reports must not be sent. The ERS flag is normally set on ICMP messages and other network layer error reports. It does not suppress the normal response to ICMP queries or similar network layer queries (CNLP echo request). If both the RFD and ERS flags are set, the fragmentation report is sent. (This definition allows a larger range of possibilities than simply over-riding the RFD flag would; a sender not desiring this behavior can see to it that RFD is clear.)

Time To Live: The time to live is an 8-bit count, nominally in seconds. Each hop is required to decrement TTL by at least one. A hop that holds a datagram for an unusual amount of time (more than 2 seconds, a typical example being a wait for a subnetwork connection establishment) should subtract the entire waiting time in seconds (rounded upward) from the TTL.

Forward Cache Identifier: Each datagram carries a 32-bit field, called "forward cache identifier," that is updated (if the information is available) at each hop. This field's value is derived from ICMP messages sent back by the next hop router, a routing protocol (e.g., RAP), or some other method. The FCI is used to expedite routing decisions by preserving knowledge where possible between consecutive routers. It can also be used to make datagrams stay within reserved flows, circuits, and mobile host tunnels. If an FCI is not available, this field must be zero, the SAO and DAO flags must be clear, and both destination and source addresses must appear in the datagram.

Datagram Length: The 32-bit length of the entire datagram in octets. A datagram can therefore be up to 4,294,967,295 bytes in overall length. Particular networks normally impose lower limits.

Transport Protocol: The transport layer protocol. For example, TCP is 6.

Checksum: The checksum is a 16-bit checksum of the entire header, using the familiar algorithm used in IP version 4.

Destination: The destination address, a count byte followed by the destination NSAP with the zero selector omitted. This field is present only if the DAO flag is zero. If the count field is not 3 modulo 4 (the destination is not an integral multiple of 32-bit words) zero bytes are added to pad to the next multiple of 32 bits. These pad bytes are not required to be ignored: routers may rely on them being zero.

Source: The source address, in the same format as the destination. Present only if the SAO flag is zero. The source is padded in the same way as destination to arrive at a 32-bit boundary.

Options: Options may follow. They are variable length, and always 32-bit aligned. If the MRO flag in the header is not set, routers will usually not look at or take action on any option, regardless of the setting of the class field.

Multicasting

The multicast-enable option permits multicast forwarding of the CATNIP datagram on subnetworks that directly support media layer multicasting. This is a vanishing species, even in 10Mbps Ethernet, given the increasing prevalence of switching hubs. It also (perhaps more usefully) permits a router to forward the datagram on multiple paths when a multicast routing algorithm has established such paths. There is no option data.

Note that there is no special address space for multicasting in CATNIP. Multicast destination addresses can be allocated anywhere by any administration or authority. This supports a number of differing models of addressing. It does require that the transport layer protocol know that the destination is multicast; this is desirable in any case. (For example, the transport layer will probably want to set the ERS flag.)

continued on next page

CATNIP *(continued)*

On an IEEE 802.x (ISO 8802.x) type media, the last 23 bits of the address (not including the 0 selector) are used in combination with the multicast group address assigned to the Internet to form the media address when forwarding a datagram with the multicast enable option from a router to an attached network, provided that the datagram was not received on that network with either multicast or broadcast media addressing. A host may send a multicast datagram either to the media multicast address (the IP catenet model) or media unicast to a router which is expected to repeat it to the multicast address within the entire level I area or to repeat copies to the appropriate end systems within the area on non-broadcast media (the more general CLNP model.)

Network Layer translation

The translating host or router must reassemble datagrams that have been fragmented before translation. Where the translation is being done by the destination host (for example, the case of a native CATNIP host receiving IP version 4 datagrams), this is similar to the present fragmentation model.

When it is being done by an intermediate router (acting as an internetwork layer gateway) the router should use all of source, destination, and datagram ID for identification of fragments. Note that destination is used implicitly in the usual reassembly at the destination. If the fragments take different paths through the net, and arrive at different translation points, the datagram is lost.

The objective of translation is to be able to upgrade systems, both hosts and routers, in whatever order desired by their owners. Organizations must be able to upgrade any given system without reconfiguration or modification of any other, and existing hosts must be able to interoperate essentially forever. (Non-CATNIP routers will probably be effectively eliminated at some point, except where they exist in their own remote or isolated corners.)

Each CATNIP system, whether host or router, must be able to recognize adjacent systems in the topology that are (only) IP version 4, CLNP, or IPX and call the appropriate translation routine just before sending the datagram.

OSI CLNP

The translation between CLNP and the CATNIP compressed form of the datagrams is the simplest case for CATNIP, since the addresses are the same and need not be extended. The resulting CATNIP datagrams may omit the source and destination addresses as explained previously, and may be mixed with uncompressed datagrams on the same subnetwork link. Alternatively, a subnetwork may operate entirely in CATNIP, converting all transit traffic to CATNIP datagrams, even if FCIs that would make the compression effective are not available.

Similarly, all network datagram formats with CATNIP mappings may be compressed into the common form, providing a uniform transit network service, with common routing protocols (such as IS–IS).

Internet Protocol

All existing version 4 numbers are defined as belonging to the Internet by using a new AFI, to be assigned to IANA by ISO. This document uses 192 at present for clarity in examples; it is to be replaced with the assigned AFI. The AFI specifies that the IDI is two bytes long, containing an administrative domain number.

The AD (Administrative Domain), identifies an administration which may be an international authority (such as the existing InterNIC), a national administration, or a large multi-organization (e.g., a government). The idea is that there should not be more than a few hundred of these at first, and eventually thousands or tens of thousands at most.

AD numbers are assigned by IANA. Initially, the only assignment is the number 0.0, assigned to the InterNIC, encompassing the entire existing version 4 Internet.

The mapping from/to version 4 IP addresses is as follows:

Length	AFI	IDI...	DSP...
7	192	AD Number	Version 4 Address

While the address (DSP) is initially always the 4-byte, version 4 address, it can be extended to arbitrary levels of subnetting within the existing Internet numbering plan. Hosts with DSPs longer than 4 bytes will not be able to interoperate with version 4 hosts.

Novell IPX

The *Internetwork Packet Exchange* (IPX) protocol, developed by Novell based on the XNS protocol (Xerox Network System) has many of the same capabilities as the Internet and OSI protocols. At first look, it appears to confuse the network and transport layers, as IPX includes both the network layer service and the user datagram service of the transport layer, while SPX (Sequenced Packet Exchange) includes the IPX network layer and provides service similar to TCP or TP4. This turns out to be mostly a matter of the naming and ordering of fields in the packets, rather than any architectural difference.

IPX uses a 32-bit LAN network number, implicitly concatenated with the 48-bit MAC layer address to form an Internet address. Initially, the network numbers were not assigned by any central authority, and thus were not useful for inter-organizational traffic without substantial prior arrangement. There is now an authority established by Novell to assign unique 32-bit numbers and blocks of numbers to organizations that desire inter-organization networking with the IPX protocol.

The Novell/IPX numbering plan uses an ICD, to be assigned, to designate an address as an IPX address. This means Novell uses the authority (AFI=47)(ICD=Novell) and delegates assignments of the following 32 bits.

An IPX address in the common form looks like:

Length	AFI	IDI...	DSP...
13	47 (hex)	Novell ICD	Network and MAC Address

This will always be followed by two bytes of zero padding when it appears in a common network layer datagram. Note that the socket numbers included in the native form IPX address are part of the transport layer.

continued on next page

CATNIP *(continued)*

SIPP

It may seem a little odd to describe the interaction with SIPP (version 6 of IP) which is now only another experimental candidate for the next generation of network layer protocols. However, if SIPP is deployed, whether or not as the protocol of choice for replacement of IP version 4, there will then be four network protocols to accommodate. It is prudent to investigate how SIPP could then be integrated into the common addressing plan and datagram format.

SIPP defines 64 bit addresses, which are included in the NSAP addressing plan under the Internet AFI as AD number 0.1. It is not clear at this time what administration will hold the authority for the SIPP numbering plan.

Length	AFI	IDI...	DSP...
11	192	AD (0.1)	SIPP 64-bit Address

The SIPP addressing method (the definition of the 64 bits) will not be described here. We note only that in the cases in which SIPP is intended to interoperate directly with IP version 4, the last 32 bits of the address is the IP version 4 address. This permits a convenient set of translations without disturbing the transport protocols.

The SIPP proposal also includes an encapsulated-tunnel mechanism called IPAE, to address some of the issues that are designed into CATNIP. The CATNIP direct translation does not use the SIPP–IPAE packet formats. IPAE also specifies a "mapping table" for prefixes. This table is kept up-to-date by periodic FTP transfers from a "central site." The CATNIP definitions leave the problem of prefix selection when converting into SIPP firmly within the scope of the SIPP–IPAE proposal, and possible methods are not described here.

In translating from SIPP (IPv6) to CATNIP (IPv7), the only unusual aspect is that SIPP defines some things that are normally considered options to be "payloads" overloaded onto the transport protocol numbering space. Fortunately, the only one that need be considered is fragmentation; a fragmented SIPP datagram must be reassembled prior to conversion. Other "payloads" such as routing are ignored (translated verbatim) and will normally simply fail to achieve the desired effect.

Translation to SIPP is simple, except for the difficult problem of inventing the "prefix" if an implementation wants to support translating Internet AD 0.0 numbers into the SIPP addressing domain.

Internet DNS

CATNIP addresses are represented in the DNS with the NSAP RR. The data in the resource record is the NSAP, including the zero selector at the end. The zone file syntax for the data is a string of hexadecimal digits, with a period "." inserted between any two octets where desired for readability. For example:

The inverse (PTR) zone is `.NSAP`, with the CATNIP address (reversed). That is, like `.IN-ADDR.ARPA`, but with `.NSAP` instead. The octets are represented as hexadecimal numbers, with leading 0's. (Zero is always written as ".00.") This respects the difference in actual authority: the IANA is the authority for the entire space rooted in `.IN-ADDR.ARPA`. in the version 4 Internet, while in the new Internet it holds the authority only for `C0.NSAP`.

The domain `00.00.C0.NSAP` is to be delegated by IANA to the InterNIC. (Understanding that in present practice the InterNIC is the operator of the authoritative root.)

Security

The CATNIP design permits the direct use of the present proposals for network layer security being developed in the IPSEC WG of the IETF. There are a number of detailed requirements; the most relevant being that network layer datagram translation must not affect (cannot affect) the transport layers, since the TPDU is mostly inaccessible to the router. For example, the translation into IPX will only work if the port numbers are shadowed into the plaintext security header.

References

[1] A. Lyman Chapin, David M. Piscitello, *Open Systems Networking*, Addison-Wesley, Reading, Massachusetts, 1993.

[2] Radia Perlman, *Interconnections: Bridges and Routers,* Addison-Wesley, Reading, Massachusetts, 1992.

[3] Jon Postel, editor, "Internet Protocol," RFC 791, September 1981.

[4] Jon Postel, editor, "Internet Control Message Protocol," RFC 792, September 1981.

[5] Jon Postel, editor, "Transmission Control Protocol," RFC 793, September 1981.

[6] Jon Postel, "NCP/TCP transition plan," RFC 801, November 1981.

[7] J. Mogul & S. Deering, "Path MTU Discovery," RFC 1191, November 1990.

[8] D. Provan, "Tunneling IPX Traffic through IP Networks," RFC 1234, June 1991.

[9] J. Moy, "OSPF Version 2," RFC 1247, July 1991.

[10] D. Clark, L. Chapin, V. Cerf, R. Braden, R. Hobby, "Towards the Future Internet Architecture," RFC 1287 December 1991.

[11] Z. Wang, J. Crowcroft, "Two-tier address structure for the Internet: A solution to the problem of address space exhaustion," RFC 1335, May 1992.

[12] V. Fuller, T. Li, J. Yu, K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[13] R. W. Callon, "TCP and UDP with Bigger Addresses (TUBA), A simple proposal for Internet addressing and routing," RFC 1347, June 1992.

[14] E. Gerich, "Guidelines for Management of IP Address Space," RFC 1466, May 1993.

[15] Robert Ullmann, "TP/IX: The Next Internet," RFC 1475, June 1993.

[16] Robert Ullmann, "RAP: Internet Route Access Protocol," RFC 1476, June 1993.

[17] D. Piscitello, "Use of ISO CLNP in TUBA Environments," RFC 1561, December 1993.

[18] Crocker, D., "The ROAD to a New IP," *ConneXions,* Volume 6, No. 11, November 1992.

ROBERT ULLMANN is currently doing network protocol design for Lotus Development Corporation in Cambridge, MA. Robert is a famous alligator wrestler, roué, philanthropic medical researcher, and metallic alloy. He speaks several extra-terrestrial languages and repairs acorns for relaxation. His antigravity demonstrations in the early 1970's brought him fame (some might say notoriety) and a large buy-out arrangement by various airlines and government organizations. He spends Tuesdays quietly, at home, polishing his goldfish. E-mail: `rullmann@crd.lotus.com`

MICHAEL McGOVERN is currently Manager of Documentation for Software House, Inc. in Waltham, MA. Michael has worked in high-tech as a computer operator, software engineer, system manager, and technical writer. He eats only berries, MYLAR™, and tektites; his other personal habits are not a matter of public record. E-mail: `scrivner@world.std.com`

TUBA: CLNP as IPng

by
Peter S. Ford, Los Alamos National Laboratory,
Yakov Rekhter, IBM, T. J. Watson Research Center,
Mark Knopper, Ameritech Advanced Data Services,
Richard Colella, NIST

Introduction

The Internet services a large and diverse community of users. Mark Lottor's latest survey of hosts on the Internet yielded a count of 2.3 million attached systems. Continuing growth of the Internet is leading to the following four problems:

- Exhaustion of the available class B network numbers;
- Exhaustion of all network numbers (class A, B, and C);
- Overhead associated with the volume of routing information, and;
- Total IP address space exhaustion.

The Internet community has developed a *Classless Inter-Domain Routing* architecture (CIDR) to address the first three problems in the context of the current Internet Protocol (IPv4) [7] [6]. CIDR exploits hierarchical routing and makes better use of the existing 32 bit IPv4 address space. This will enable the use of the current generation of IP through the end of this century [1].

The *IP: Next Generation* (IPng) efforts are focused on solving the problem of IPv4 address exhaustion. The problem of IPv4 address exhaustion cannot be solved using 32-bit addresses. TUBA (TCP and UDP with Bigger Addresses) is a solution for IPng [4] [18].

TUBA has several design objectives:

- Addressing and scalable routing capabilities for arbitrarily large Internets.
- Support of multiprotocol Internets. The TUBA effort believes that IPng should be capable of facilitating convergence of network layer protocols such as IPv4, Novell's IPX, and CLNP, pointing to a future where most internets converge on the use of a common underlying internet protocol.
- Minimize the risk and cost of transitioning to IPng, given the Internet's large operational infrastructure and user community.
- Provide a base for continued evolution of Internet services.

Description

TUBA has two components: a replacement for the network layer protocol and a transition strategy. These components are essentially orthogonal; the protocols and transition strategies in the various IPng proposals can be mixed and matched, as desired (e.g., the transition strategy proposed for TUBA could be used in conjunction with SIPP as the network layer protocol). The choices made for TUBA specifically address the design objectives above.

TUBA Network Layer

TUBA replaces the Internet network layer, IPv4, with the OSI *Connectionless Network Layer Protocol,* CLNP [13] [9]. TUBA systems allow the current Internet applications, such as *Telnet, FTP, Mosaic* and electronic mail, to operate using CLNP as the network layer protocol. CLNP shares most of the architectural features and functionality of IPv4, but is distinguished by its use of flexible, variable length addresses called *Network Service Access Points* (NSAPs).

The IETF OSI operations planning group has developed an NSAP addressing plan [2] that meets the objective of addressing for an Internet of the size anticipated in the next century. The flexibility of NSAPs allows the embedding of other network layer addresses, such as IPv4 and Novell IPX. Other organizations have selected NSAPs for addressing, such as the ATM Forum's use of NSAPs for addressing within ATM networks.

CLNP is already understood by most developers of Internet products and operators of the Internet infrastructure. Implementations of CLNP exist today for most host and router systems. The use of the CLNS MIB [24] allows management of CLNP systems using SNMP. CLNP has been widely deployed in wide area networks as the map below illustrates:

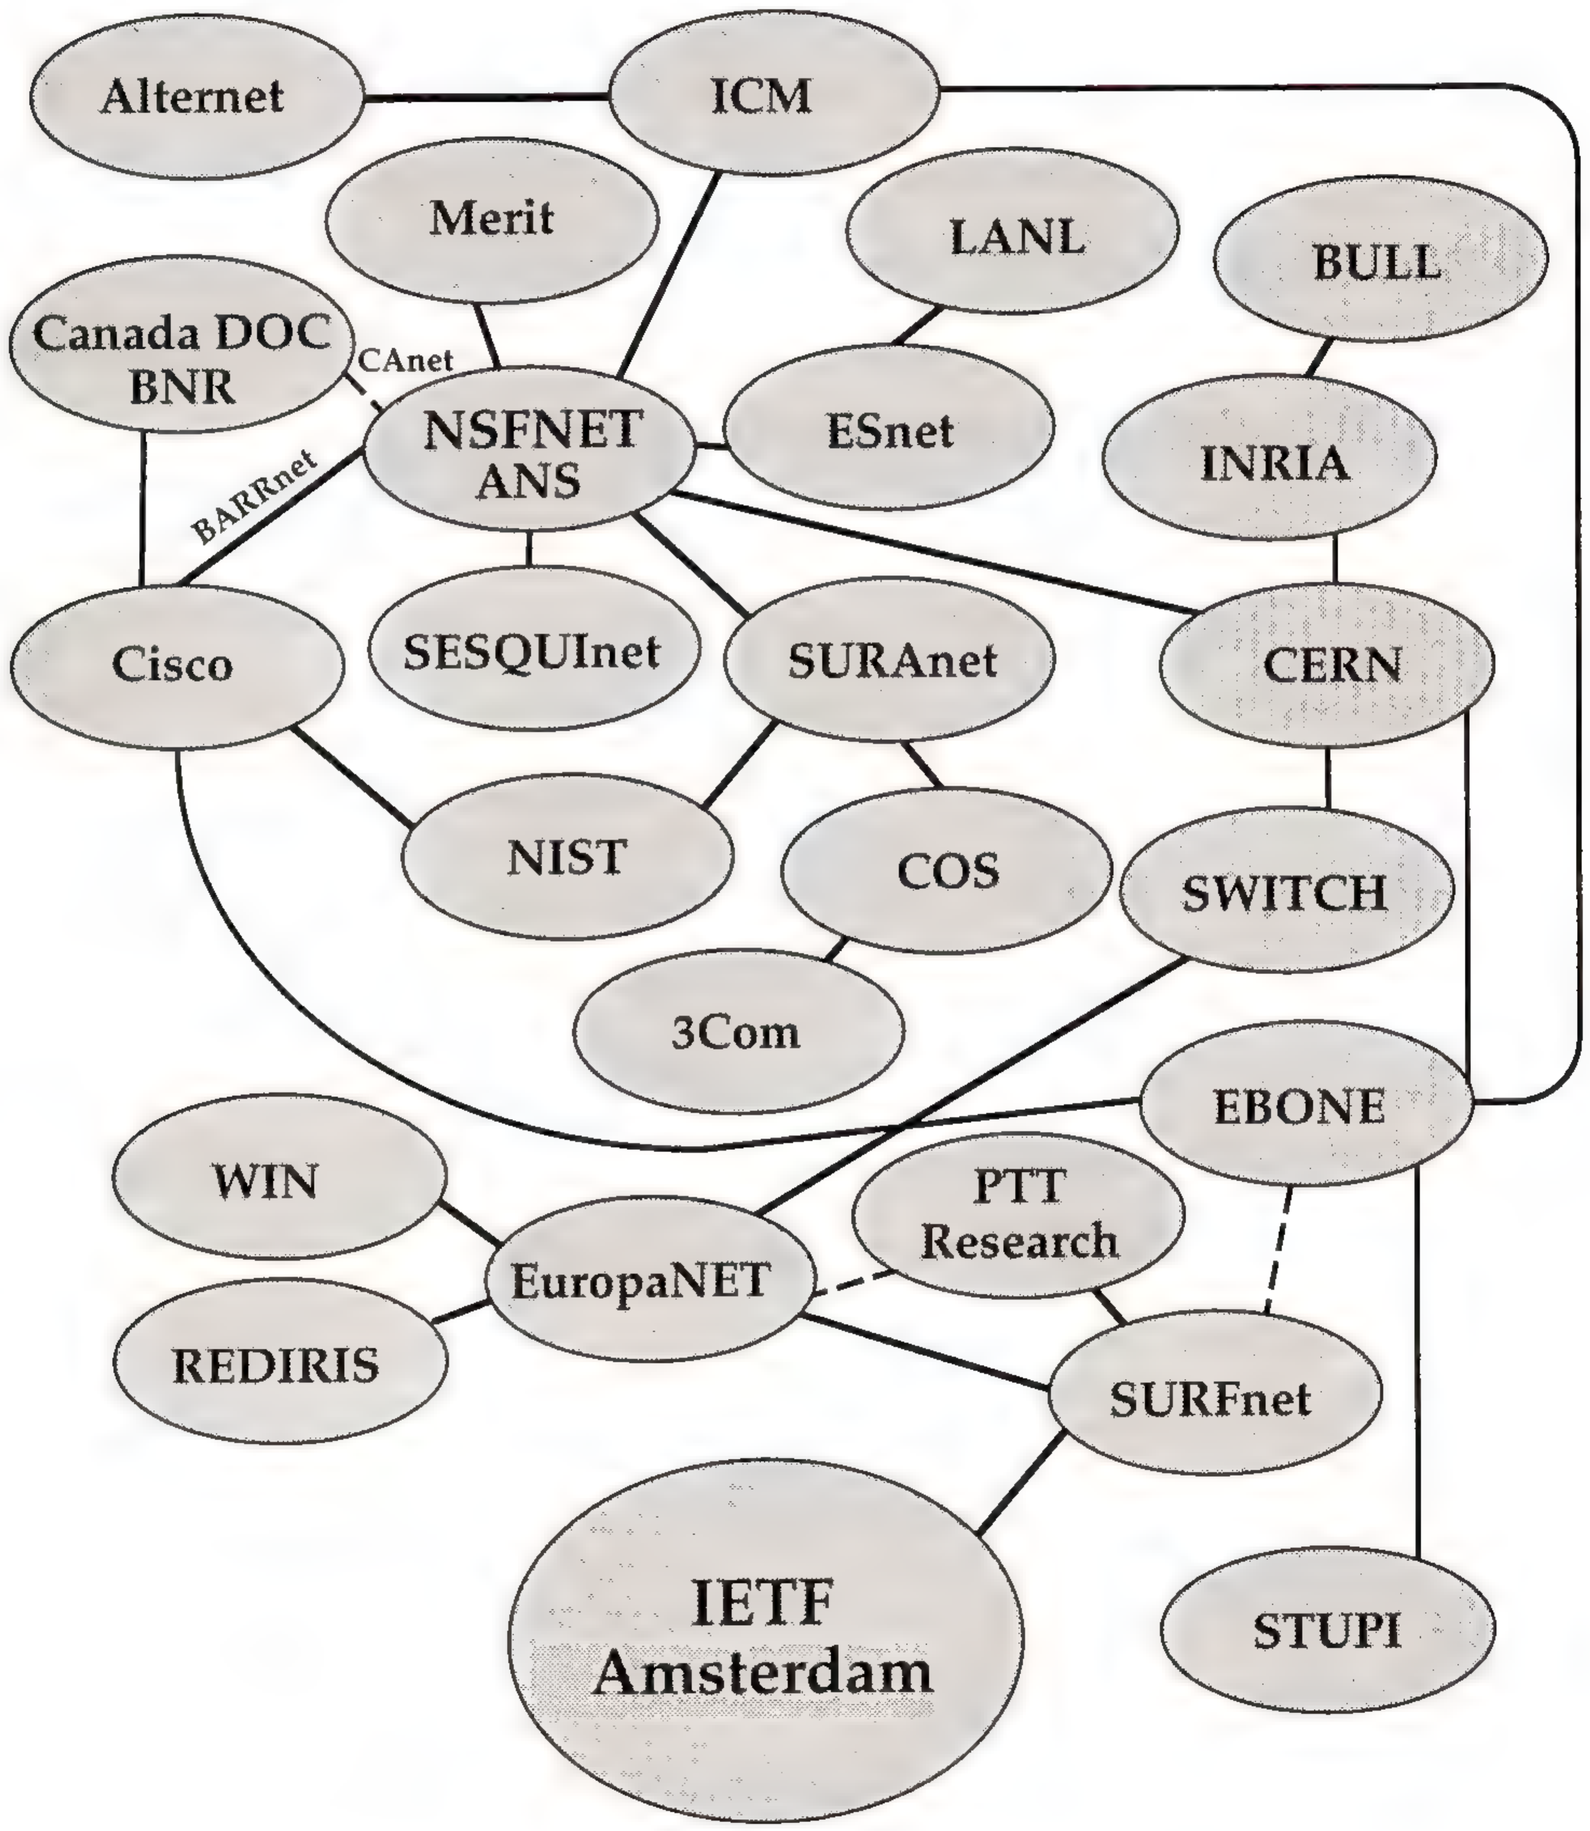

Figure 1: July 1993 TUBA Infrastructure and Participants

One major advantage of using CLNP as a replacement for IPv4 is that the routing architecture has already been specified, standardized, and implemented in products. Routing protocols used with CLNP include IDRP [16] [12], IS–IS [15] [25], and ES–IS [14] [10]. The TUBA routing architecture and protocols (IDRP and IS–IS) follow IPv4's CIDR architecture and protocols (BGP and OSPF). IDRP and IS–IS are sufficiently flexible that they can be used to route multiple network layer protocols including IPv4, IPX and SIPP.

continued on next page

TUBA: CLNP as IPng *(continued)*

As the customer base of the Internet grows and the demand for new services continues, it is anticipated that new functionality in the Internet suite of protocols will be introduced, e.g., security [8], multicasting [20], autoconfiguration [17], mobility, resource reservation, and enhanced support for policy-based routing. Many of these problems can and are being addressed within the IPv4 context and can be analogously implemented in TUBA systems. For several of these functions, there are standardization efforts and experimental implementations in progress for CLNP. Other functions, like autoconfiguration in the ES–IS protocol, have been largely completed. The TUBA working group is tracking the efforts of the CDPD cellular data group in the use of CLNP in mobile internets. Efforts in the IETF *Source Demand Routing Protocol* (SDRP) Working Group [5] can be applied to CLNP to solve policy routing issues such as provider selection.

Transition

The TUBA transition plan uses a *Dual Stack* (DS) strategy. The Dual Stack transition strategy assumes the existence of a multiprotocol infrastructure supporting both IPv4 and CLNP, and the development of Dual Stack hosts that concurrently use IPv4 and CLNP. The current Internet infrastructure is well positioned for a TUBA transition since a large number of the current Internet service providers already support IPv4 and CLNP. The TUBA effort has successfully demonstrated *Telnet* and the *File Transfer Protocol* (FTP) running on top of CLNP between sites located in the U.S., Europe, and Canada.

The Internet uses the *Domain Name System* (DNS) [21] to map Internet names to IPv4 addresses. The TUBA effort has specified how the DNS represents the mapping from Internet names to NSAPs and the reverse mapping from NSAPs to names. Support for these mappings has been implemented in the most recent version of BIND and several sites on the Internet report both IPv4 and NSAP addresses in response to DNS queries.

In the DS transition plan, software in new and updated hosts supports Internet transport simultaneously on top of both IPv4 and CLNP. When a DS host is attached to the DS infrastructure it is configured to use both IPv4 and CLNP. This fact is advertised to other hosts on the Internet by putting both its IPv4 address and its NSAP into the DNS.

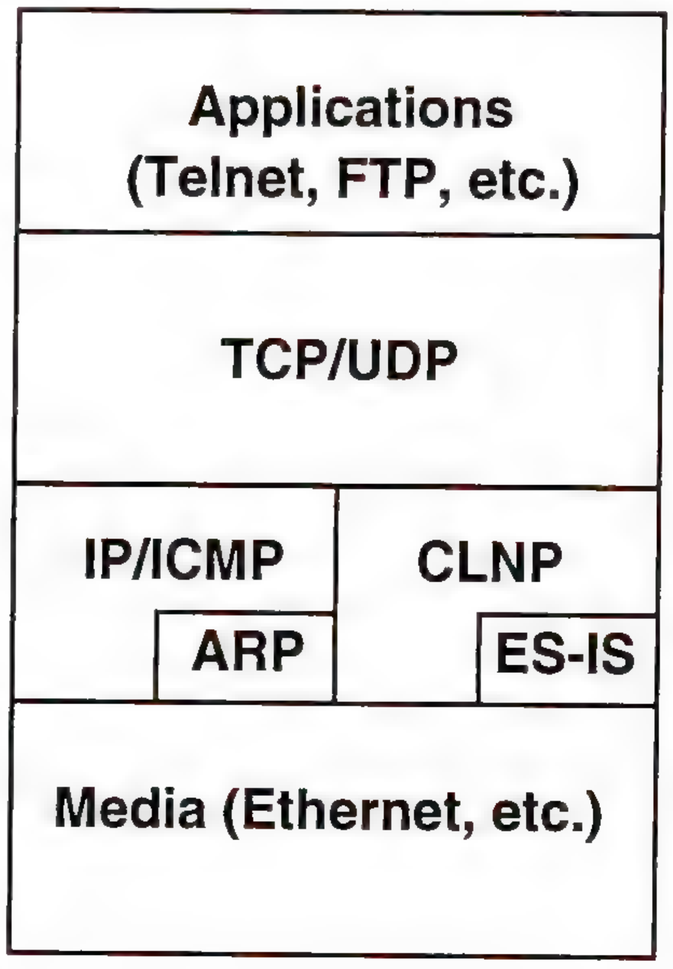

Figure 2: Dual Stack for Hosts (CLNP and IP)

DS hosts recognize each other by the existence of their NSAPs in the DNS. When attempting communication, a DS host queries the DNS for both addresses of the target host. If the DNS returns both an IPv4 and a CLNP address, then the target is another DS host connected to the DS infrastructure, and the communication can take place using CLNP. If only an IPv4 address is returned, then the target is not a DS host (or is, but has not yet been put onto the DS infrastructure) and IPv4 is used for communication. This strategy allows a host-by-host transition to TUBA thus simplifying the transition.

Implementation

Implementing TUBA in host systems is relatively simple given existing CLNP and Internet transport implementations; for example, a Berkeley UNIX implementation of TUBA required on the order of 100 lines of UNIX kernel changes. Similar ease in implementing prototypes of TUBA functionality has been reported by Cisco and 3Com on top of their existing CLNP router software base. An MS-DOS implementation, which involved development of CLNP as well as TUBA, was produced in under two months. The size of a CLNP implementation, including a substantial amount of debugging code, occupies on the order of 10K bytes in BSDi UNIX kernel space.

Application programs that make use of the Internet protocol suite need to be generalized to operate with both IPv4 and CLNP protocols. These changes are usually fairly simple, requiring small changes to algorithms and data structures that manipulate network layer addresses. For example, the File Transfer Protocol (FTP) needs to be extended to be able to use longer network layer addresses which are sent as part of the FTP PORT command [22]. Enhancing the API for Internet services to handle TUBA DNS requests will greatly simplify these efforts. Any IPng will require similar changes to be made to support hosts once the 32 bit IPv4 address space is exhausted.

Interoperation

There are implementations of TUBA for BSDi UNIX, MS-DOS (*Telnet* and *FTP*), and Sun. Cisco and 3Com have prototype versions of their router software which support TUBA *Telnet*. These systems have been successfully tested for interoperability using *Telnet* and *FTP*. Using the existing CLNP infrastructure, international interoperation has been demonstrated. During Summer 1993, TUBA *telnets* transited five internets switching CLNP to intercommunicate between the IETF meeting hall in Amsterdam, The Netherlands, and TUBA hosts in New Mexico, USA. There is an encapsulation of CLNP on top of IPv4 that allows networks not connected to a CLNP-capable Internet services provider to tunnel CLNP over IPv4 only infrastructure to get world wide CLNP interconnectivity.

Summary

Use of CLNP and the TUBA routing architecture provides addresses and scalable routing capabilities for Internets of practically unlimited size. This addresses the IPv4 address space exhaustion problem.

By using an existing network layer and a simple dual stack transition strategy, TUBA provides a cost-effective and low-risk solution for IPng. Further, by relying on CLNP and associated routing protocols, TUBA preserves the investment in the deployed routing infrastructure and technology.

Success of the Internet is predicated on its ability to gracefully introduce new functionality and services. CLNP, and the use of NSAP addressing, provides a flexible framework for the introduction of new functionality such as security, mobility, and multicasting.

continued on next page

TUBA: CLNP as IPng *(continued)*

More information

To join the TUBA mailing list send mail to `tuba-request@merit.edu`. Information on the TUBA effort can be obtained via anonymous FTP from the directory: `merit.edu:pub/tuba-archive` or from the *Gopher* server `explorer.nsap.research.ptt.nl`. Documentation of CLNP and related routing protocols are in `merit.edu:pub/iso`.

References

[1] Braun, Hans-Werner, Ford, Peter S., Rekhter, Yakov, "CIDR and the Evolution of the Internet Protocol," *ConneXions,* Volume 7, No. 9, September 1993.

[2] Colella, Richard, Gardner, Ella, and Callon, Ross, "Guidelines for OSI NSAP Allocation in the Internet," RFC 1237, July 1991.

[3] Callon, Ross, "Use of OSI IS–IS for Routing in TCP/IP and Dual Environments," RFC 1195, December 1990.

[4] Callon, Ross, "TCP And UDP with Bigger Addresses (TUBA)," RFC 1347, June 1992.

[5] Estrin, Deborah, Rekhter, Yakov, Hotz, "A Unified Approach to Inter-Domain Routing," RFC 1322, May 1992.

[6] Ford, Peter S., Braun, Hans-Werner, Rekhter, Yakov, "Improving the Routing and Addressing of IP," *IEEE Network,* May 1993.

[7] Fuller, V., Li, T., Yu, J., Varadhan, K., "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[8] Glenn, K. Robert, "Integrated Network Layer Security Protocol," Internet Draft, work in progress, September 1993.

[9] Hagens, Rob, "Components of OSI: CLNP," *ConneXions,* Volume 3, No. 10, October 1989.

[10] Hagens, Rob, "Components of OSI: ES–IS Routing," *ConneXions,* Volume 3, No. 8, August 1989.

[11] Hares, S., Wittbrodt, C., "Essential Tools for the OSI Internet," Internet Draft, work in progress, March 1993.

[12] Hares, Sue, "Components of OSI: Inter Domain Routing Protocol," *ConneXions,* Volume 6, No. 5, 1992.

[13] ISO/IEC, "Protocol for Providing the Connectionless-mode Network Service," International Standard 8473, 1986.

[14] ISO/IEC, "End System to Intermediate System Routeing Exchange Protocol for use in Conjunction with the Protocol for the Provision of the Connectionless-mode Network Service," International Standard 9542, 1987.

[15] ISO/IEC, "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473)," International Standard 10589, 1992.

[16] ISO/IEC, "Protocol for Exchange of Inter-Domain Routeing Information among Intermediate Systems to support Forwarding of ISO 8473 PDUs," International Standard 10747, 1993.

[17] Katz, Dave, "Dynamic System ID Assignment," (work in progress).

[18] Katz, Dave, Ford, Peter S., "TUBA: Replacing IP with CLNP," *IEEE Network,* Volume 7, No. 3, May 1993.

[19] Manning, B., Colella, R., "DNS NSAP Resource Records," Internet Draft, work in progress, December 1993.

[20] Marlow, Dave, "Host Group Extensions for CLNP Multicasting," (work in progress).

[21] Mockapetris, Paul V., "Domain names—concepts and facilities," RFC 1034, November 1987.

[22] Piscitello, Dave, "FTP Operation Over Big Address Records (FOOBAR)," RFC 1545, November 1993.

[23] Piscitello, Dave, "Use of ISO CLNP in TUBA Environments," RFC 1561, December 1993.

[24] Satz, G., "CLNS MIB for use with ISO 8473 CLNP and ISO 9542 ES–IS," RFC 1238, July 1991.

[25] Tsuchiya, Paul, "Components of OSI: IS–IS Intra-Domain Routing," *ConneXions,* Volume 3, No. 8, August 1989.

[26] Comer, D. E., *Internetworking with TCP/IP, Volume I, Principles, Protocols, and Architecture,* Second Edition, ISBN 0-13-468505-9, Prentice Hall, 1991.

[27] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., "Address Allocation for Private Internets," RFC 1597, March 1994.

[28] Crocker, D., "The ROAD to a New IP," *ConneXions,* Volume 6, No. 11, November 1992.

[29] B. Carpenter, "IPng White Paper on transition and other considerations," Internet Draft, work in progress, March 1994.

[30] F. Kastenholz, C. Partridge, "Technical Criteria for Choosing IP: The Next Generation (IPng)," Internet Draft, work in progress, March 1994.

PETER SEWALL FORD is a member of the technical staff at Los Alamos National Laboratory where he is currently working with the U.S. National Science Foundation on the NSFNET and developing software to manage and instrument an 800 Mbit/sec crossbar network. Mr. Ford has also worked at Los Alamos' Center for Nonlinear Studies, on UNIX systems for the Computer Science Department at the University of Utah, and on distributed database systems and languages at Britton-Lee, Inc. His university degree is in General Studies from the University of Michigan. E-mail: `peter@goshawk.lanl.gov`

YAKOV REKHTER holds M.S. in Physics from St. Petersburg University, Russia, M.S. in Computer Science from New York University, and Ph.D. in Computer Science from Polytechnic University. He has been with IBM since 1984, and is currently a Research Staff Member and a manager of High-Performance Communication group at T. J. Watson Research Center. Dr. Rekhter was one of the leading architects, as well as the major software developer, of the NSFNET Backbone Phase II. Present activities include work on Classless Inter-Domain Routing (CIDR), work on the Unified Approach to Inter-Domain Routing (under the contract with the National Science Foundation), work on supporting host mobility, and work on supporting IP over Fibre Channel subnetworks. His Internet e-mail address is: `yakov@watson.ibm.com`

MARK KNOPPER has recently moved to Ameritech Advanced Data Services, where he is director of Network Information Infrastructure. Previously he was at Merit Network, Inc. since 1980, and served as manager of Internet Engineering for the NSFNET backbone services. Mark is co-chair of the TUBA working group, and is a member of the IPng Directorate. E-mail: `mak@merit.edu`

RICHARD COLELLA has been with the U.S. National Institute of Standards and Technology (NIST) as a Computer Scientist since 1984. During the last two years, Mr. Colella has been working towards the addition of OSI's CLNP to the multi-protocol Internet. Most recently in this area he has been working on various aspects of the TUBA proposal, including the development of a public domain prototype. His Internet efforts also included chairing an IETF working group that produced a guide for the allocation of CLNP addresses in the Internet. His Internet e-mail address is: `colella@nist.gov`

Simple Internet Protocol Plus (SIPP) Overview

by Robert M. Hinden, Sun Microsystems

Introduction

This article presents an overview of the *Simple Internet Protocol Plus* (SIPP) which is one of the candidates being considered in the *Internet Engineering Task Force* (IETF) for the next version of the Internet Protocol (the current version is usually referred to as IPv4). This article is not intended to be a detailed presentation of all of the features and motivation for SIPP, but is intended to give the reader an overview of the proposal. It is also not intended that this be an implementation specification, but given the simplicity of the central core of SIPP, an implementor familiar with IPv4 could probably construct a basic working SIPP implementation from reading this overview.

SIPP is a new version of IP which is designed to be an evolutionary step from IPv4. It is a natural increment to IPv4. It can be installed as a normal software upgrade in internet devices and is interoperable with the current IPv4. Its deployment strategy was designed to not have any "flag" days. SIPP is designed to run well on high performance networks (e.g., ATM) and at the same time is still efficient for low bandwidth networks (e.g., wireless). In addition, it provides a platform for new internet functionality that will be required in the near future.

This article describes the work of the IETF SIPP working group. Several individuals deserve specific recognition. These include Steve Deering, Paul Francis, Dave Crocker, Bob Gilligan, Bill Simpson, Ran Atkinson, Bill Fink, Erik Nordmark, Christian Huitema, Sue Thompson, and Ramesh Govindan.

Key issues for IPng

There are several key issues that should be used in the evaluation of any next generation internet protocol. Some are very straightforward. For example the new protocol must be able to support large global internetworks. Others are less obvious. There must be a clear way to transition the current installed base of IP systems. It doesn't matter how good a new protocol is if there isn't a practical way to transition the current operational systems running IPv4 to the new protocol.

Growth

Growth is the basic issue which caused there to be a need for a next generation IP. If anything is to be learned from our experience with IPv4 it is that the addressing and routing must be capable of handling reasonable scenarios of future growth. It is important that we have an understanding of the past growth and where the future growth will come from.

Currently IPv4 serves what could be called the computer market. The computer market has been the driver of the growth of the Internet. It comprises the current Internet and countless other smaller internets which are not connected to the Internet. Its focus is to connect computers together in the large business, government, and university education markets. This market has been growing at an exponential rate. One measure of this is that the number of networks in current Internet (23,494 as of 1/28/94) is doubling approximately every 12 months. The computers which are used at the endpoints of Internet communications range from PCs to Supercomputers. Most are attached to Local Area Networks (LANs) and the vast majority are not mobile.

New markets

The next phase of growth will probably not be driven by the computer market. While the computer market will continue to grow at significant rates due to expansion into other areas such as schools (elementary through high school) and small businesses, it is doubtful it will continue to grow at an exponential rate. What is likely to happen is that other kinds of markets will develop. These markets will fall into several areas. They all have the characteristic that they are extremely large. They also bring with them a new set of requirements which were not as evident in the early stages of IPv4 deployment. The new markets are also likely to happen in parallel with each other. It may turn out that we will look back on the last ten years of Internet growth as the time when the Internet was small and only doubling every year. The challenge for an IPng is to provide a solution which solves today's problems and is attractive in these emerging markets.

Nomadic personal computing devices seem certain to become ubiquitous as their prices drop and their capabilities increase. A key capability is that they will be networked. Unlike the majority of today's networked computers they will support a variety of types of network attachments. When disconnected they will use RF wireless networks, when used in networked facilities they will use infrared attachment, and when docked they will use physical wires. This makes them an ideal candidate for internetworking technology as they will need a common protocol which can work over a variety of physical networks. These types of devices will become consumer devices and will replace the current generation of cellular phones, pagers, and personal digital assistants. In addition to the obvious requirement of an internet protocol which can support large scale routing and addressing, they will require an internet protocol which imposes a low overhead and supports auto configuration and mobility as a basic element. The nature of nomadic computing requires an internet protocol to have built in authentication and confidentiality. It also goes without saying that these devices will need to communicate with the current generation of computers. The requirement for low overhead comes from the wireless media. Unlike LANs which will be very high speed, the wireless media will be several orders of magnitude slower due to constraints on available frequencies, spectrum allocation, and power consumption.

Another market is networked entertainment. The first signs of this emerging market are the proposals being discussed for 500 channels of television, video on demand, etc. This is clearly a consumer market. The possibility is that every television set will become an Internet host. As the world of digital high definition television approaches, the differences between a computer and a television will diminish. As in the previous market, this market will require an Internet protocol which supports large scale routing and addressing, and auto configuration. This market also requires a protocol suite which imposes the minimum overhead to get the job done. Cost will be the major factor in the selection of a technology to use.

Another market which could use the next generation IP is device control. This consists of the control of everyday devices such as lighting equipment, heating and cooling equipment, motors, and other types of equipment which are currently controlled via analog switches and in aggregate consume considerable amounts of power. The potential cost savings to save power by networking these devices is very large. The size of this market is enormous and requires solutions which are simple, robust, easy to use, and very low cost.

continued on next page

Simple Internet Protocol Plus *(continued)*

The challenge for the IETF in the selection of an IPng is to pick a protocol which meets today's requirements and also matches the requirements of these emerging markets. These markets will happen with or without an IETF IPng. If the IETF IPng is a good match for these new markets it is likely to be used. If not, these markets will develop something else. They will not wait for an IETF solution. If this should happen it is probable that because of the size and scale of the new markets the IETF Internet protocols would be supplanted. If the IETF IPng is not appropriate for use in these markets, it is also probable that they will each develop their own protocols, perhaps proprietary. These new protocols would not interoperate with each other. The opportunity for the IETF is to select an IPng which has a reasonable chance to be used in these emerging markets. This would have the very desirable outcome of creating an immense, interoperable, world-wide information infrastructure created with open protocols. The alternative is a world of disjoint networks with protocols controlled by individual vendors.

Transition

At some point in the next three to seven years the Internet will require a deployed new version of the Internet protocol. Two factors are driving this: *routing* and *addressing*. Global Internet routing based on the on 32-bit addresses of IPv4 is becoming increasingly strained. IPv4 address do not provide enough flexibility to construct efficient hierarchies which can be aggregated. The deployment of *Classless Inter-Domain Routing* (CIDR) [1] is extending the life time of IPv4 routing by a number of years, the effort to manage the routing will continue to increase. Even if the IPv4 routing can be scaled to support a full IPv4 Internet, the Internet will eventually run out of network numbers. There is no question that an IPng is needed, but only a question of when.

The challenge for an IPng is for its transition to be complete before IPv4 routing and addressing break. The transition will be much easier if IPv4 address are still globally unique. The two transition requirements which are the most important are flexibility of deployment and the ability for IPv4 hosts to communicate with IPng hosts. There will be IPng-only hosts, just as there will be IPv4-only hosts. The capability must exist for IPng-only hosts to communicate with IPv4-only hosts globally while IPv4 addresses are globally unique.

The deployment strategy for an IPng must be as flexible as possible. The Internet is too large for any kind of controlled rollout to be successful. The importance of flexibility in an IPng and the need for interoperability between IPv4 and IPng was well stated in a message to the SIPP mailing list by Bill Fink, who is responsible for a portion of NASA's operational internet. In his message he said:

> "Being a network manager and thereby representing the interests of a significant number of users, from my perspective it's safe to say that the transition and interoperation aspects of any IPng is *the* key first element, without which any other significant advantages won't be able to be integrated into the user's network environment. I also don't think it wise to think of the transition as just a painful phase we'll have to endure en route to a pure IPng environment, since the transition/coexistence period undoubtedly will last at least a decade and may very well continue for the entire lifetime of IPng, until it's replaced with IPngng and a new transition. I might wish it was otherwise but I fear they are facts of life given the immense installed base.

"Given this situation, and the reality that it won't be feasible to coordinate all the infrastructure changes even at the national and regional levels, it is imperative that the transition capabilities support the ability to deploy the IPng in the piecemeal fashion...with no requirement to need to coordinate local changes with other changes elsewhere in the Internet...

"I realize that support for the transition and coexistence capabilities may be a major part of the IPng effort and may cause some headaches for the designers and developers, but I think it is a duty that can't be shirked and the necessary price that must be paid to provide as seamless an environment as possible to the end user and his basic network services such as e-mail, ftp, Gopher, X-Window clients, etc...

"The bottom line for me is that we must have interoperability during the extended transition period for the base IPv4 functionality..."

Another way to think about the requirement for compatibility with IPv4 is to look at other product areas. In the product world, backwards compatibility is very important. Vendors who do not provide backward compatibility for their customers usually find they do not have many customers left. For example, chip makers put considerable effort into making sure that new versions of their processor always run all of the software that ran on the previous model. It is unlikely that Intel would develop a new processor in the X86 family that did not run DOS and the tens of thousands of applications which run on the current versions of X86 processors.

Operating system vendors go to great lengths to make sure new versions of their operating systems are binary compatible with their old version. For example the labels on most PC or MAC software usually indicate that they require OS version XX or greater. It would be foolish for Microsoft come out with a new version of Windows which did not run the applications which ran on the previous version. Microsoft even provides the ability for windows applications to run on their new OS called NT. This is an important feature. They understand that it was very important to make sure that the applications which run on Windows also run on NT.

The same requirement is also true for IPng. The Internet has a large installed base. Features need to be designed into an IPng to make the transition as easy as possible. As with processors and operating systems, it must be backwards compatible with IPv4. Other protocols have tried to replace TCP/IP, for example XTP and OSI. One element in their failure to reach widespread acceptance was that neither had any transition strategy other than running in parallel (sometimes called dual stack). New features alone are not adequate to motivate users to deploy new protocols. IPng must have a great transition strategy and new features.

History of the SIPP effort

The SIPP working group represents the evolution of three different IETF working groups focused on developing an IPng. The first was called *IP Address Encapsulation* (IPAE) and was chaired by Dave Crocker and Robert Hinden. It proposed extensions to IPv4 which would carry larger addresses. Much of its work was focused on developing transition mechanisms.

Somewhat later Steve Deering proposed a new protocol evolved from IPv4 called the *Simple Internet Protocol* (SIP). A working group was formed to work on this proposal which was chaired by Steve Deering and Christian Huitema.

continued on next page

Simple Internet Protocol Plus *(continued)*

SIP had 64-bit addresses, a simplified header, and options in separate extension headers. After lengthy interaction between the two working groups and the realization that IPAE and SIP had a number of common elements and the transition mechanisms developed for IPAE would apply to SIP, the groups decided to merge and concentrate their efforts. The chairs of the new SIP working group were Steve Deering and Robert Hinden.

In parallel to SIP, Paul Francis (formerly Paul Tsuchiya) had founded a working group to develop the *"P" Internet Protocol* (Pip). Pip was a new internet protocol based on a new architecture. The motivation behind Pip was that the opportunity for introducing a new internet protocol does not come very often, and given that opportunity important new features should be introduced. Pip supported variable length addressing in 16-bit units, separation of addresses from identifiers, support for provider selection, mobility, and efficient forwarding. It included a transition scheme similar to IPAE.

After considerable discussion among the leaders of the Pip and SIP working groups, they came to realize that the advanced features in Pip could be accomplished in SIP without changing the base SIP protocol as well as keeping the IPAE transition mechanisms. In essence it was possible to keep the best features of each protocol. Based on this the groups decided to merge their efforts. The new protocol was called *Simple Internet Protocol Plus* (SIPP). The chairs of the merged working group are Steve Deering, Paul Francis, and Robert Hinden.

SIPP overview

SIPP is a new version of the Internet Protocol, designed as a successor to IP version 4 [8]. SIPP is assigned IP version number 6. SIPP was designed to take an evolutionary step from IPv4. It was not a design goal to take a radical step away from IPv4. Functions which work in IPv4 were kept in SIPP. Functions which didn't work were removed. The changes from IPv4 to SIPP fall primarily into the following categories:

- *Expanded Routing and Addressing Capabilities:* SIPP increases the IP address size from 32 bits to 64 bits, to support more levels of addressing hierarchy and a much greater number of addressable nodes. SIPP addressing can be further extended, in units of 64 bits, by a facility equivalent to IPv4's Loose Source and Record Route option, in combination with a new address type called "cluster addresses" which identify topological regions rather than individual nodes. The scalability of multicast routing is improved by adding a "scope" field to multicast addresses.

- *Header Format Simplification:* Some IPv4 header fields have been dropped or made optional, to reduce the common-case processing cost of packet handling and to keep the bandwidth cost of the SIPP header almost as low as that of IPv4, despite the increased size of the addresses. The basic SIPP header is only four bytes longer than IPv4.

- *Improved Support for Options:* Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new options in the future.

- *Quality-of-Service Capabilities:* A new capability is added to enable the labeling of packets belonging to particular traffic "flows" for which the sender requests special handling, such as non-default quality of service or "real-time" service.

- *Authentication and Privacy Capabilities:* SIPP includes the definition of extensions which provide support for authentication, data integrity, and confidentiality. This is included as a basic element of SIPP.

Header format

The SIPP protocol consists of two parts, the basic SIPP header and SIPP Options. The header is shown below:

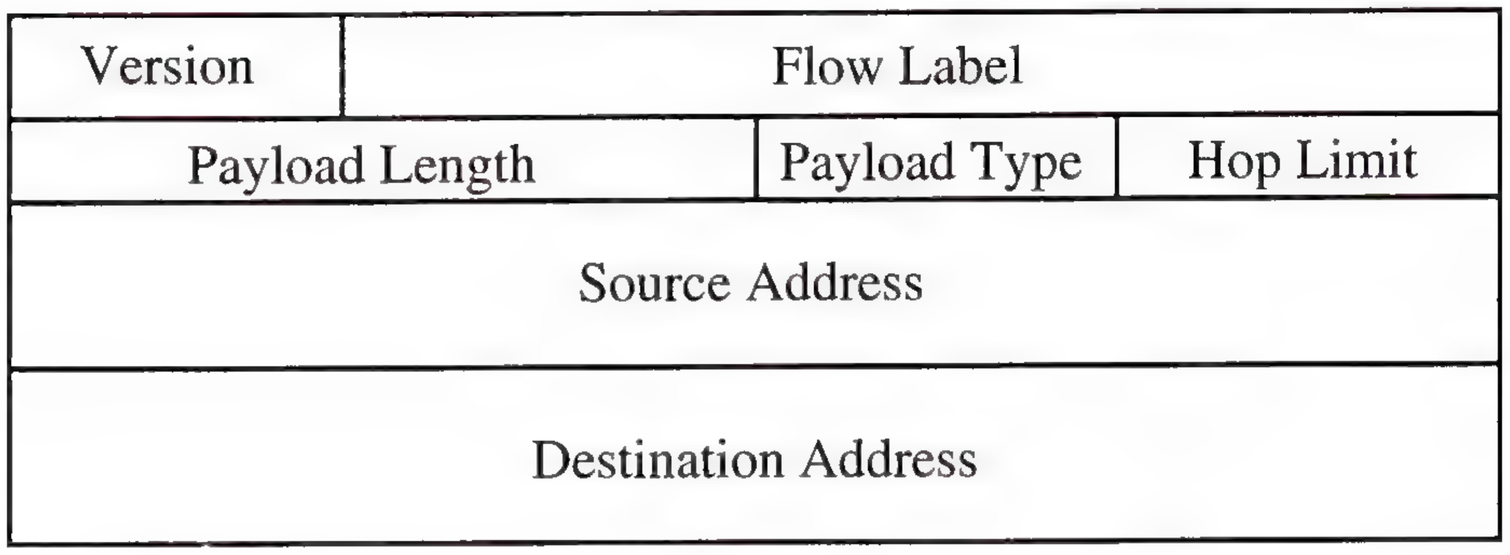

Version: 4-bit Internet Protocol version number = 6.

Flow Label: 28-bit field. See SIPP Quality of Service section.

Payload Length: 16-bit unsigned integer. Length of payload, i.e., the rest of the packet following the SIPP header, in octets.

Payload Type: 8-bit selector. Identifies the type of header immediately following the SIPP header. Uses the same values as the IPv4 Protocol field [13].

Hop Limit: 8-bit unsigned integer. Decremented by 1 by each node that forwards the packet. The packet is discarded if Hop Limit is decremented to zero.

Source Address: 64 bits. An address of the initial sender of the packet. See [4] for details.

Destination Address: 64 bits. An address of the intended recipient of the packet (possibly not the ultimate recipient, if an optional *Routing Header* is present).

Options

SIPP includes an improved option mechanism over IPv4. SIPP options are placed in separate headers that are located between the SIPP header and the transport-layer header in a packet. Most SIPP option headers are not examined or processed by any router along a packet's delivery path until it arrives at its final destination. This facilitates a major improvement in router performance for packets containing options. In IPv4 the presence of any options requires the router to examine all options. The other improvement is that unlike IPv4, SIPP options can be of arbitrary length and the total amount of options carried in a packet is not limited to 40 bytes. This feature plus the manner in which they are processed, permits SIPP options to be used for functions which were not practical in IPv4. A good example of this is the SIPP Authentication and Security Encapsulation options.

In order to improve the performance when handling subsequent option headers and the transport protocol which follows, SIPP options are always an integer multiple of 8 octets long, in order to retain this alignment for subsequent headers.

continued on next page

Simple Internet Protocol Plus *(continued)*

The SIPP option headers which are currently defined are:

Option	Function
Routing	Extended Routing (like IPv4 loose source route)
Fragmentation	Fragmentation and Reassembly
Authentication	Integrity and Authentication
Security Encapsulation	Confidentiality
Hop-by-Hop Option	Special options which require hop by hop processing

Addressing

SIPP addresses are 64-bits long and are identifiers for individual nodes and sets of nodes. There are three types of SIPP addresses. These are *unicast, cluster,* and *multicast.* Unicast addresses identify a single node. Cluster addresses identify a group of nodes, that share a common address prefix, such that a packet sent to a cluster address will be delivered to one member of the group. Multicast addresses identify a group of nodes, such that a packet sent to a multicast address is delivered to all of the nodes in the group.

SIPP supports addresses which are twice the number of bits as IPv4 addresses. These addresses support an address space which is four billion (2^{32}) times the size of IPv4 addresses (2^{32}). Another way to say this is that SIPP supports four billion internets each the size of the maximum IPv4 Internet. That is enough to allow each person on the planet to have their own internet. Even with several layers of hierarchy (with assignment utilization similar to IPv4) this would allow for each person on the planet to have their own internet each holding several thousand hosts.

In addition, SIPP supports extended addresses using the routing option. This capability allows the address space to grow to 128-bits, 192-bits (or even larger) while still keeping the address units in manageable 64-bit units. This permits the addresses to grow while keeping the routing algorithms efficient because they continue to operate using 64-bit units.

Unicast addresses

There are several forms of unicast address assignment in SIPP. These are global hierarchical unicast addresses, local-use addresses, and IPv4-only host addresses.

Global unicast addresses

Global unicast addresses are used for global communication. They are the most common SIPP address and are similar in function to IPv4 addresses. Their format is:

1	n bits	m bits	p bits	63–n–m–p
C	PROVIDER ID	SUBSCRIBER ID	SUBNET ID	NODE ID

The first bit is the IPv4 *compatibility bit,* or C-bit. It indicates whether the node represented by the address is IPv4 or SIPP. SIPP addresses are provider-oriented. That is, the high-order part of the address is assigned to Internet service providers, which then assign portions of the address space to subscribers, etc. This usage is similar to assignment of IP addresses under CIDR. The SUBSCRIBER ID distinguishes among multiple subscribers attached to the provider identified by the PROVIDER ID. The SUBNET ID identifies a topologically connected group of nodes within the subscriber network identified by the subscriber prefix. The NODE ID identifies a single node among the group of nodes identified by the subnet prefix.

Local-use address

A local-use address is a unicast address that has only local routability scope (within the subnet or within a subscriber network), and may have local or global uniqueness scope. They are intended for use inside of a site for "plug and play" local communication, for bootstrapping up to a single global addresses, and as part of an address sequence for global communication. Their format is:

The NODE ID is a identifier which much be unique in the domain in which it is being used. In most cases these will use a node's IEEE-802 address. The SUBNET ID identifies a specific subnet in a site. This permits a large private internet to be constructed without any other address allocation.

Local-use addresses have two primary benefits. First, for sites or organizations that are not (yet) connected to the global Internet, there is no need to request an address prefix from the global Internet address space. Local-use addresses can be used instead. If the organization connects to the global Internet, it must then form addresses with global routability scope.

The second benefit of local-use addresses is that they can hold much larger NODE IDs, which makes possible a very simple form of autoconfiguration of addresses. In particular, a node may discover a SUBNET ID by listening to a Router Advertisement messages on its attached link(s), and then fabricating a SIPP address for itself by using its link-level address as the NODE ID on that subnet.

An auto-configured local-use address may be used by a node as its own identification for communication within the local domain, possibly including communication with a local address server to obtain a global SIPP address. The details of host auto-configuration are described in [6].

IPv4-only addresses

SIPP unicast addresses are assigned to IPv4-only hosts as part of the IPAE scheme for transition from IPv4 to SIPP. Such addresses have the following form:

The highest-order bit of a SIPP address is called the IPv4 compatibility bit or the C bit. A C bit value of 1 identifies an address as belonging to an IPv4-only node.

The IPv4 node's 32-bit IPv4 address is carried in the low-order 32 bits of the SIPP address. The remaining 31 bits are used to carry HIGHER-ORDER SIPP PREFIX, such as a service-provider ID.

Cluster addresses

Cluster addresses are unicast addresses that are used to reach the "nearest" one (according to unicast routing's notion of nearest) of the set of boundary routers of a cluster of nodes identified by a common prefix in the SIPP unicast routing hierarchy. These are used to identify a set of nodes. The cluster address, when used as part of an address sequence, permits a node to select which of several providers it wants to carry its traffic. In this example there would be a cluster address for each provider. This capability is sometimes called "source selected policies."

continued on next page

Simple Internet Protocol Plus *(continued)*

Cluster addresses have the general form:

n bits	64–n bits
Cluster prefix	0000000000000000000000000000000000

Multicast addresses

A SIPP multicast address is an identifier for a group of nodes. A node may belong to any number of multicast groups. Multicast addresses have the following format:

1	7	4	4	48 bits
C	1111111	FLGS	SCOP	Group ID

Where:

C = IPv4 compatibility bit.

1111111 in the rest of the first octet identifies the address as being a multicast address.

FLGS is a set of 4 flags:

0	0	0	T

The high-order 3 flags are reserved, and must be initialized to 0.

T = 0 indicates a permanently-assigned ("well-known") multicast address, assigned by the global internet numbering authority.

T = 1 indicates a non-permanently-assigned ("transient") multicast address.

SCOP is a 4-bit multicast scope value used to limit the scope of the multicast group. The values are:

0	reserved	8	intra-organization scope
1	intra-node scope	9	(unassigned)
2	intra-link scope	10	(unassigned)
3	(unassigned)	11	intra-community scope
4	(unassigned)	12	(unassigned)
5	intra-site scope	13	(unassigned)
6	(unassigned)	14	global scope
7	(unassigned)	15	reserved

GROUP ID identifies the multicast group, either permanent or transient, within the given scope.

Routing

Routing in SIPP is almost identical to IPv4 routing under CIDR except that the addresses are 64-bit SIPP addresses instead of 32-bit IPv4 addresses. This is true even when extended addresses are being used. With very straightforward extensions, all of IPv4's routing algorithms (OSPF, BGP, RIP, IDRP, etc.) can used to route SIPP.

SIPP also includes simple routing extensions which support powerful new routing functionality. These capabilities include:

- Provider Selection (based on policy, performance, cost, etc.)
- Host Mobility (route to current location)
- Auto-Readdressing (route to new address)
- Extended Addressing (route to "sub-cloud")

The new routing functionality is obtained by creating sequences of SIPP addresses using the SIPP Routing option. The routing option is used by a SIPP source to list one or more intermediate nodes (or topological clusters) to be "visited" on the way to a packet's destination. This function is very similar in function to IPv4's Loose Source and Record Route option.

The identification of a specific transport connection is done by only using the first (source) and last (destination) address in the sequence. This permits the middle addresses in the address sequence to change (in the cases of mobility, provider changes, site readdressing, etc.) without disrupting the transport connection.

In order to make address sequences a general function, SIPP hosts are required to reverse routes in a packet it receives containing address sequences in order to return the packet to its originator. This approach is taken to make SIPP host implementations from the start support the handling and reversal of source routes. This is the key for allowing them to work with hosts which implement the new features such as provider selection or extended addresses.

Three examples show how the extended addressing can be used. In these examples, address sequences are shown by a list of individual addresses separated by commas. For example:

SRC, I1, I2, I3, DST

Where the first address is the source address, the last address is the destination address, and the middle addresses are intermediate addresses.

For these examples assume that two hosts, H1 and H2 wish to communicate. Assume that H1 and H2's sites are both connected to providers P1 and P2. A third wireless provider, PR, is connected to both providers P1 and P2.

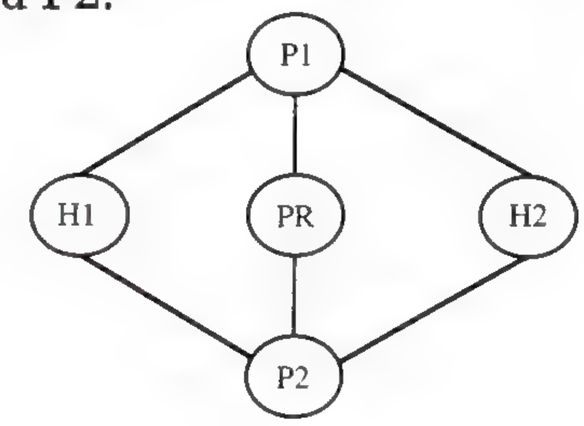

The simplest case (no use of address sequences) is when H1 wants to send a packet to H2 containing the addresses:

H1, H2

When H2 replied it would reverse the addresses and construct a packet containing the addresses:

H2, H1

In this example either provider could be used, and H1 and H2 would not be able to select which provider traffic would be sent to and received from.

If H1 decides that it wants to enforce a policy that all communication to/from H2 can only use provider P1, it would construct a packet containing the address sequence:

H1, P1, H2

This ensures that when H2 replies to H1, it will reverse the route and the reply would also travel over P1. The addresses in H2's reply would look like:

H2, P1, H1

continued on next page

Simple Internet Protocol Plus *(continued)*

If H1 became mobile and moved to provider PR, it could maintain (not breaking any transport connections) communication with H2, by sending packets that contain the address sequence:

H1, PR, P1, H2

This would ensure that when H2 replied it would enforce H1's policy of exclusive use of provider P1 and send the packet to H1's new location on provider PR. The reversed address sequence would be:

H2, P1, PR, H1

The address extension facility of SIPP can be used for provider selection, mobility, readdressing, and extended addressing. It is a simple but powerful capability.

Quality-of-service capabilities

The Flow Label field in the SIPP header may be used by a host to label those packets for which it requests special handling by SIPP routers, such as non-default quality of service or "real-time" service. This labeling is important in order to support applications which require some degree of consistent throughput, delay, and/or jitter. The Flow Label is a 28-bit field, internally structured into three subfields as follows:

R (Reserved): 1-bit subfield. Initialized to zero for transmission; Ignored on reception.

DP (Drop Priority): 3-bit unsigned integer. Specifies the priority of the packet, relative to other packets from the same source, for being discarded by a router under conditions of congestion. Larger values indicates a greater willingness by the sender to allow the packet to be discarded.

Flow ID: 24-bit subfield used to identify a specific flow.

A flow is a sequence of packets sent from a particular source to a particular (unicast or multicast) destination for which the source desires special handling by the intervening routers. There may be multiple active flows from a source to a destination, as well as traffic that is not associated with any flow. A flow is identified by the combination of a Source Address and a non-zero Flow ID. Packets that do not belong to a flow carry a Flow ID of zero.

A Flow ID is assigned to a flow by the flow's source node. New Flow IDs must be chosen (pseudo-)randomly and uniformly from the range 1 to FFFFFF hex. The purpose of the random allocation is to make any set of bits within the Flow ID suitable for use as a hash key by the routers, for looking up the special-handling state associated with the flow. A Flow ID must not be re-used by a source for a new flow while any state associated with the previous usage still exists in any router.

The Drop Priority subfield provides a means separate from the Flow ID for distinguishing among packets from the same source, to allow a source to specify which of its packets are to be discarded in preference to others when a router cannot forward them all. This is useful for applications like video where it is preferable to drop packets carrying screen updates rather than the packets carrying the video synchronization information.

Security

The current Internet has a number of security problems and lacks effective privacy and authentication mechanisms below the application layer. SIPP remedies these shortcomings by having two integrated options that provide security services. These two options may be used singly or together to provide differing levels of security to different users. This is very important because different user communities have different security needs.

The first mechanism, called the "SIPP Authentication Header," is an option which provides authentication and integrity (without confidentiality) to SIPP datagrams. While the option is algorithm-independent and will support many different authentication techniques, the use of keyed MD5 is proposed to help ensure interoperability within the worldwide Internet. This can be used to eliminate a significant class of network attacks, including host masquerading attacks. The use of the SIPP Authentication Header is particularly important when source routing is used with SIPP because of the known risks in IP source routing. Its placement at the internet layer can help provide host origin authentication to those upper layer protocols and services that currently lack meaningful protections. This mechanism should be exportable by vendors in the United States and other countries with similar export restrictions because it only provides authentication and integrity, and specifically does not provide confidentiality. The exportability of the SIPP Authentication Header encourages its widespread implementation and use.

The second security option provided with SIPP is the "SIPP Encapsulating Security Header." This mechanism provides integrity and confidentiality to SIPP datagrams. It is simpler than some similar security protocols (e.g., SP3D, ISO NLSP) but remains flexible and algorithm-independent. To achieve interoperability within the global Internet, the use of DES CBC is proposed as the standard algorithm for use with the SIPP Encapsulating Security Header.

Transition mechanisms

The two key motivations in the SIPP transition mechanisms are to provide direct interoperability between IPv4 and SIPP hosts and to allow the user population to adopt SIPP in an a highly diffuse fashion. The transition must be incremental, with few or no critical interdependencies, if it is to succeed. The SIPP transition allows the users to upgrade their hosts to SIPP, and the network operators to deploy SIPP in routers, with a minimum of coordination between the two.

The mechanisms and policies of the SIPP transition are called "IPAE." Having a separate term serves to highlight those features designed specifically for transition. Once an acronym for an encapsulation technique to facilitate transition, the term "IPAE" now is mostly historical.

The IPAE transition is based on five key elements:

- A 64-bit SIPP addressing plan that encompasses the existing 32-bit IPv4 addressing plan. The 64-bit plan will be used to assign addresses for both SIPP and IPv4 nodes at the beginning of the transition. Existing IPv4 nodes will not need to change their addresses, and IPv4 hosts being upgraded to SIPP keep their existing IPv4 addresses as the low-order 32 bits of their SIPP addresses. Since the SIPP addressing plan is a superset of the existing IPv4 plan, SIPP hosts are assigned only a single 64-bit address, which can be used to communicate with both SIPP and IPv4 hosts.

continued on next page

Simple Internet Protocol Plus *(continued)*

- A mechanism for encapsulating SIPP traffic within IPv4 packets so that the IPv4 infrastructure can be leveraged early in the transition. Most of the "SIPP within IPv4 tunnels" can be automatically configured.
- Algorithms in SIPP hosts that allow them to directly interoperate with IPv4 hosts located on the same subnet and elsewhere in the Internet.
- A mechanism for translating between IPv4 and SIPP headers to allow SIPP-only hosts to communicate with IPv4-only hosts and to facilitate IPv4 hosts communicating over over a SIPP-only backbone.
- An optional mechanism for mapping IPv4 addresses to SIPP address to allow improved scaling of IPv4 routing. At the present time given the success of CIDR, this does not look like it will be needed in a transition to SIPP. If Internet growth should continue beyond what CIDR can handle, it is available as an optional mechanism.

IPAE ensures that SIPP hosts can interoperate with IPv4 hosts anywhere in the Internet up until the time when IPv4 addresses run out, and afterward allows SIPP and IPv4 hosts within a limited scope to interoperate indefinitely. This feature protects for a very long time the huge investment users have made in IPv4. Hosts that need only a limited connectivity range (e.g., printers) need never be upgraded to SIPP. This feature also allows SIPP-only hosts to interoperate with IPv4-only hosts.

The incremental upgrade features of IPAE allow the host and router vendors to integrate SIPP into their product lines at their own pace, and allows the end users and network operators to deploy SIPP on their own schedules.

The interoperability between SIPP and IPv4 provided by IPAE also has the benefit of extending the lifetime of IPv4 hosts. Given the large installed base of IPv4, changes to IPv4 in hosts are nearly impossible. Once an IPng is chosen, most of the new feature development will be done on IPng. New features in IPng will increase the incentives to adopt and deploy it.

Why SIPP?

There are a number of reasons why SIPP should be selected as the IETF's IPng. It solves the Internet scaling problem, provides a flexible transition mechanism for the current Internet, and was designed to meet the needs of new markets such as nomadic personal computing devices, networked entertainment, and device control. It does this in a evolutionary way which reduces the risk of architectural problems.

Ease of transition is a key point in the design of SIPP. It is not something was was added in at the end. SIPP is designed to interoperate with IPv4. Specific mechanisms (C-bit, embedded IPv4 addresses, etc.) were built into SIPP to support transition and compatibility with IPv4. It was designed to permit a gradual and piecemeal deployment without any dependencies.

SIPP supports large hierarchical addresses which will allow the Internet to continue to grow and provide new routing capabilities not built into IPv4. It has cluster addresses which can be used for policy route selection and has scoped multicast addresses which provide improved scalability over IPv4 multicast. It also has local use addresses which provide the ability for "plug and play" installation.

SIPP is designed to have performance better than IPv4 and work well in low bandwidth applications like wireless. Its headers are less expensive to process than IPv4 and its 64-bit addresses are chosen to be well matched to the new generation of 64-bit processors. Its compact header minimizes bandwidth overhead which makes it ideal for wireless use.

SIPP provides a platform for new Internet functionality. This includes support for real-time flows, provider selection, host mobility, end-to-end security, auto-configuration, and auto-reconfiguration.

In summary, SIPP is a new version of IP. It can be installed as a normal software upgrade in Internet devices. It is interoperable with the current IPv4. Its deployment strategy was designed to not have any "flag" days. SIPP is designed to run well on high performance networks (e.g., ATM) and at the same time is still efficient for low bandwidth networks (e.g., wireless). In addition, it provides a platform for new Internet functionality that will be required in the near future.

Status of the SIPP effort

There are many active participants in the SIPP working group. Groups making active contributions include:

Group	*Activity*
Beame & Whiteside	Implementation (PC)
Bellcore	Implementation (SunOS), DNS, ICMP specs.
Digital Equipment Corp.	Implementation (Alpha/OSF, Open VMS)
INRIA	Implementation (BSD, BIND), DNS, OSPF specs.
INESC	Implementation (BSD/Mach/x-kernel)
InterCon	Implementation (MAC)
MCI	Phone Conferences
Merit	IDRP for SIPP Specification
Naval Research Lab	Implementation (BSD) Security Design
Network General	Implementation (Sniffer)
SGI	Implementation (IRIX, NetVisualizer)
Sun	Implementation (Solaris 2.x, Snoop)
TGV	Implementation (VMS)
Xerox PARC	Protocol Design
Bill Simpson	Implementation (KA9Q)

As of the time this article was written there were a number of SIPP and IPAE implementations. These include:

Implementation	*Status*
BSD/Mach	Completed (*telnet,* NFS, AFS, UDP)
BSD/Net/2	In Progress
BIND	Code done
DOS & Windows	Completed (*telnet, ftp, tftp, ping*)
IRIX	In progress (*ping*)
KA9Q	In progress (*ping,* TCP)
Mac OS	Completed (*telnet, ftp, finger, ping*)
NetVisualizer	Completed (SIP & IPAE)
Open VMS	In Progress
OSF	In Progress
Sniffer	Completed (SIP & IPAE)
Snoop	Completed (SIP & IPAE)
Solaris	Completed (*telnet, ftp, tftp, ping*)
Sun OS	In Progress
VMS	Completed (*telnet, ftp*)

continued on next page

Simple Internet Protocol Plus *(continued)*

Additional information

The documentation listed in the reference sections can be found in one of the IETF Internet Draft directories or in the archive site for the SIPP working group. This is located at:

`ftp.parc.xerox.com` in the `/pub/SIPP` directory.

In addition other material relating to SIPP (such as *PostScript* versions of presentations on SIPP) can also be found in the SIPP working group archive.

A *Mosaic* page has been created for the SIPP working group and installed on the Internet. It can be found at:

`http://town.hall.org` under the `new protocol` heading.

To join the SIPP working group, send electronic mail to:

`sipp-request@sunroof.eng.sun.com`

An archive of mail sent to this mailing list can be found in the IETF directories at `cnri.reston.va.us`.

References

[1] Fuller, V., Li, T., Yu, J., Varadhan, K., "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[2] S. Deering, "Simple Internet Protocol Plus (SIPP) Specification," Internet Draft, work in progress, February 1994.

[3] R. Gilligan, et al, "IPAE: The SIPP Interoperability & Transition Mechanism," Internet Draft, work in progress, November 1993.

[4] S. Deering, et al, "Simple Internet Protocol Plus (SIPP): Routing and Addressing," Internet Draft, work in progress, January 1994.

[5] R. Govindan, et al, "ICMP and IGMP for SIPP Specification," Internet Draft, work in progress.

[6] W. Simpson, "SIPP Auto-configuration," Internet Draft, work in progress.

[7] P. Francis, "Simple Internet Protocol Plus (SIPP): Unicast Hierarchical Address Assignment," Internet Draft, work in progress, January 1994.

[8] J. Postel, "Internet Protocol," RFC 791, September 1981.

[9] R. Atkinson, "SIPP Authentication Payload," Internet Draft, work in progress, January 1994.

[10] R. Atkinson, "SIPP Encapsulating Security Payload (ESP)," Internet Draft, work in progress, January 1994.

[11] R. Atkinson, "SIPP Security Architecture," Internet Draft, work in progress, January, 1994.

[12] W. Simpson, "SIPP Neighbor Discovery," Internet Draft, work in progress, December 1993.

[13] J. Reynolds, J. Postel, "Assigned Numbers," RFC 1340, July 1992.

ROBERT HINDEN holds a B.S.E.E. and an M.S. in Computer Science from Union College, Schenectady, New York. He works for Sun Microsystems where he is responsible for the department which develops internet protocols for Sun's operating systems. Prior to this he worked at Bolt, Beranek, and Newman, Inc. on a variety of internetwork related projects including the first operational internet router and one of the first TCP/IP implementations. He has been active in the IETF since 1985 and served for the past seven years as the Area Director for Routing in the Internet Engineering Steering Group and is the co-chair of the SIPP working group, and has previously chaired the IP over ATM and the Open Routing working groups. He can be reached on the Internet at `hinden@eng.sun.com`.

A User's View of the Next Generation of IP (IPng)

by Eric Fleischman, Boeing Computer Services

Introduction

The activities within the Internet community to resolve the scaling problems of the Internet Protocol (IP) have been well documented by this and other publications. These articles have detailed the specific proposals which are being examined as potential replacements for the current version of IP, IP version 4 (IPv4), and its routing protocols. This quest to identify the future version of IP, together with the list of contending replacement protocol solutions, is often generically termed *IP: The Next Generation* and abbreviated "IPng."

Viewpoints

However, not many articles to date have examined the implications of IPng to the current TCP/IP user community. The goal of this article is to present a user's viewpoint on the IPng issues. This "user's view of IPng" must be presented with a large dose of humility. Humility is necessary because the various issues which pertain to the selection of a scalable IP replacement protocol are quite complex. These issues represent a multi-faceted problem domain which may be profitably examined from many pertinent viewpoints. Among these viewpoints are the following:

- Network Service Providers
- Computer/Network Research Community
- Vendors of Computer/Network Products. This community may be further subdivided into vendors of "Intermediate Systems" products and vendors of "End Systems" products.
- End Users (i.e., entities who use computers as an "overhead expense" to accomplish a non-computer-oriented goal). Several different end-user perspectives may be distinguished including private citizens, government, education, and industry.

Each of these diverse viewpoints provide valuable perspectives for evaluating the IPng problem domain. Each perspective potentially asks different questions which are highly relevant for the ultimate identification of the "best" IPng proposal. However, the sheer number of these differing viewpoints complicates the IPng evaluation process by adding additional concerns and requirements. This process is further complicated by the fact that different entities sharing the same "viewpoint" are themselves subject to different priorities due to "corporate culture" distinctives. For example, one would expect a small corporation to have a different list of priorities than a very large corporation even though both viewpoints are "industry" viewpoints. Similarly, entities perpetually on the technological "bleeding edge" will naturally view things differently than technologically conservative entities.

Because of the complexity of the larger context, this article will examine the implications of IPng from the point of view of a single Fortune 100 corporation which has heavily invested in TCP/IP technology in order to achieve its business goals. While this is merely a view from a single corporation, it is hoped that this viewpoint may prove to be generally relevant to the Internet Community as a whole.

Characteristics

The following 5 key characteristics describe our environment and are probably representative of other large TCP/IP deployments. We believe that understanding these characteristics is very important for obtaining insight into the implication of IPng to large user populations:

continued on next page

A User's View of IPng *(continued)*

• *Host Ratio:* Many corporations explicitly try to limit the number of their TCP/IP hosts that are directly accessible from the Internet. This is done for a variety of reasons (e.g., security). While the ratio of those hosts that have direct Internet access capabilities to those hosts without such capabilities will vary from company to company, ratios ranging from 1:1000 to 1:10,000 (or more) are not uncommon. In addition, not all corporations with TCP/IP deployments currently have Internet attachments. The implication of this point is that the state of the world-wide (IPv4) Internet address space only directly impacts a tiny percentage of the currently deployed TCP/IP hosts within a large corporation. This is true even if the entire population is currently using Internet-assigned addresses.

• *Router-to-Host Ratio:* Most corporations have significantly more TCP/IP hosts than they have IP routers. Ratios ranging between 100:1 to 600:1 (or more) are common. The implication of this point is that a transition approach which solely demands changes to routers is generally much less disruptive than an approach which demands changes to both routers and hosts.

• *Business Factor:* Large corporations exist to fulfill some business purpose such as the construction of airplanes, baseball bats, cars, or some other product offering. Computing is merely a tool to help automate business processes in order to more efficiently accomplish the business goals of the corporation. Automation is accomplished via *applications.* Data communications, operating systems, and computer hardware are simply the tools used by applications to accomplish their goals. Thus, users actually buy applications and not networking technologies. The central lesson of this point is that IPng will be deployed according to the applications which use it and not because it is a better technology.

• *Integration Factor:* Large corporations currently support many diverse computing environments. This diversity limits the effectiveness of a corporation's computing assets by hindering data sharing, application interoperability, "application portability," and software reusability. The net effect is stunted application life cycles and increased support costs. Data communications is but one of the domains which contribute towards this diversity. For example, The Boeing Company currently has deployed at least sixteen different protocol families within its networks (e.g., TCP/IP, SNA, DECnet, OSI, IPX/SPX, AppleTalk, XNS, etc.). Each distinct Protocol Family population potentially implies unique training, administrative, support, and infrastructure requirements. Consequently, corporate goals often exist to eliminate or merge diverse Data Communications Protocol Family deployments in order to reduce network support costs and to increase the number of devices which can communicate together (i.e., foster interoperability). This results in a basic abhorrence to the possibility of introducing "Yet Another Protocol" (YAP). Consequently, an IPng solution which introduces an entirely new set of protocols will be negatively viewed simply because its by-products are more roadblocks to interoperability coupled with more work, expense, and risk to support our computing resources and business goals.

Having said this, it should be observed that this abhorrence may be partially overcome by "extenuating circumstances" such as applications using IPng which meet critical end-user requirements or by broad (international) commercial support.

• *Inertia Factor:* There is a natural tendency to continue to use the current IP protocol (IPv4) regardless of the state of the Internet's IPv4 address space. Motivations supporting inertia include the following: existing application dependencies (including *Application Programming Interface* (API) dependencies); opposition to additional protocol complexity; budgetary constraints limiting additional hardware/software expenses; additional address management and naming services costs; transition costs; support costs; training costs; etc. As the number of our deployed TCP/IP hosts continues to grow towards the 100,000 mark, the inertial power of this population becomes increasingly strong.

However, inertia even exists with smaller populations simply because the cost to convert or upgrade the systems are not warranted. Consequently, pockets of older "Legacy System" technologies often exist in specific environments (e.g., we still have pockets of the archaic BSC protocol). The significance of this point is that unless there are significant business benefits to justify an IPng deployment, economics will oppose such a deployment. Thus, even if the forthcoming IPng protocol proves to be "the ultimate and perfect protocol," it is unrealistic to imagine that the entire IPv4 population will ever transition to IPng. This means that should we deploy IPng within our network, there will be an ongoing requirement for our internal IPng deployment to be able to communicate with our internal IPv4 community. This requirement is unlikely to go away with time.

Costs and benefits

Thus, the central, bottom-line question concerning IPng from the user perspective is: What are the benefits which will justify the expense of deploying IPng? At this time we can conceive of only four possible causes which may motivate us to deploy IPng:

Possible Cause:	*Possible Corporate Response:*
• Many Remote (external) Peers solely use IPng.	Gateway external systems only.
• Internet requires IPng usage.	Gateway external systems only.
• "Must have" products are tightly coupled with IPng (e.g., "flows" for real-time applications).	Upgrade internal corporate network to support IPng where that functionality is needed.
• Senior management directs IPng usage.	Respond appropriately.

It should explicitly be noted that the reasons which are compelling the Internet Community to create IPng (i.e., the scalability of IPv4 over the Internet) are not themselves adequate motivations for users to deploy IPng within their own private networks. That is, should IPng usage become mandated as a prerequisite for Internet usage, a probable response to this mandate would be to convert our few hosts with direct external access capabilities to become IPng-IPv4 application-layer gateways (i.e., dual stacks). This would leave the remainder of our vast internal TCP/IP deployment unchanged. Consequently, given gateways for external access, there may be little motivation for a company's internal network to support IPng.

continued on next page

A User's View of IPng *(continued)*

Motivations

We suspect that there are only two causes which will motivate users to widely deploy IPng:

(1) If IPng products add critical functionality which IPv4 can't provide (e.g., real time applications, multimedia applications, genuine (scalable) plug-and-play networking, etc.), users would be motivated to deploy IPng where that functionality is needed. However, these deployments must combat the "Integration Factor" and the "Inertia Factor" forces which have previously been described. This implies that there must be a significant business gain to justify such a deployment. While it is impossible to predict exactly how this conflict would "play out," it is reasonable to assume that IPng would probably be deployed according to an "as needed only" policy. Optimally, specific steps would be taken to protect the remainder of the network from the impact of these localized changes. Of course, should IPng become bundled with "killer applications" (i.e., applications which are extremely important to significantly many key business processes) then all bets are off: IPng will become widely deployed.

(2) Should IPng foster a convergence between Internet Standards and International Standards (i.e., OSI), this convergence could change IPng's destiny. That is, the networks of many large corporations are currently being driven by sets of strong, but contradictory, requirements: one set demanding compliance with Internet Standards and another set demanding compliance with International Standards. [Note: The following is a single example concerning why International Standards are important to large corporations. Corporations conducting a global business are subject to the regulations of those countries in which they trade. International commerce is regulated by governments, many of whom have placed restrictions upon data communications. These restrictions affect the data communications of a corporation's products as well as the data communications between corporations (i.e., business partners, customers, and suppliers). International Standards are the only certain bet to comply with world-wide commerce restrictions.]

If a means could be found to achieve greater synergy (integration/adoption) between Internet Standards and International Standards then corporate management may very well be inclined to mandate internal deployment of the merged standards and promote their external use. Optimally, such a synergy should offer the promise of reducing currently deployed protocol diversity (i.e., supports the "Integration Factor" force). Depending on the specific method by which this convergence is achieved, it may also partially offset the previously mentioned "Inertia Factor" force, especially if IPng proves to be a protocol which has already been deployed.

Consequently, mandating IPng to communicate over the Internet does not correspondingly imply the need for large corporations to generally support IPng within their networks. Thus, while the IPv4 scalability limitations are compelling reasons to identify a specific IPv4 replacement protocol for the Internet, other factors are at work within private corporate networks. These factors imply that large TCP/IP end users will have a continuing need to purchase IPv4 products even after IPng products have become generally available.

User requirements

This article began with an acknowledgment that the Internet community is composed of a variety of members who possess potentially differing viewpoints of the IPng problem domain.

We perceive that our vantage point would identify the following as critical end-user requirements for IPng:

• The IPng approach must permit users to slowly transition to IPng in a piecemeal fashion. Even if IPng becomes widely deployed, it is unrealistic to expect that users will ever transition all of the extensive IPv4 installed base to IPng. Consequently, the approach must indefinitely support corporate-internal communication between IPng hosts and IPv4 hosts regardless of the requirements of the world-wide Internet.

• The IPng approach must not hinder technological advances to be implemented (e.g., mobile hosts, multimedia applications, or real-time applications).

• The IPng approach is expected to eventually foster greater synergy (integration/adoption) between Internet Standards and International Standards (i.e., OSI). [Note: This may be accomplished in a variety of ways including having the Internet Standards adopted as International Standards or else having the International Standards adopted as Internet Standards.]

• The IPng approach should have "self-defining network" (i.e., "plug and play") capabilities. That is, large installations require device portability in which one may readily move devices within one's corporate network and have them autoconfigure, autoaddress, autoregister, etc. without explicit human administrative overhead at the new location—assuming that the security criteria of the new location have been met.

Conclusion

In summary, the key factor which will determine whether—and to what extent—IPng will be deployed by large end users is whether IPng will become an essential element for the construction of applications which are critically needed by our businesses. If IPng is bundled with applications which satisfy critical business needs, it will be deployed. If it isn't, it is of little relevance to the large end user. Regardless of what happens to IPng, the large mass of IPv4 devices will ensure that IPv4 will remain an important protocol for the foreseeable future and that continued development of IPv4 products is advisable.

References

[1] Crocker, D., "The ROAD to a New IP," *ConneXions,* Volume 6, No. 11, November 1992.

[2] Lottor, M., "Internet Growth (1981–1991)," RFC 1296, Jan. 1992.

[3] Solensky, F., "The Growing Internet," *ConneXions,* Volume 6, No. 5, May 1992.

[4] desJardins, R., "Internet 2000," *ConneXions,* Volume 6, No. 10, October 1992.

ERIC FLEISCHMAN is a Senior Principal Scientist (Network Architect) within Boeing Computer Services' Delivery Systems Architecture and Technical Planning group. He is currently participating in a number of corporate-wide technology assessment and system integration activities. Previous to this he was actively involved in designing the Boeing Enterprise Network and assisted several of Boeing's strategic computer vendors in the design of their network products and technologies. Eric has previously worked as a Software Engineer for Victor Technologies (Scotts Valley, Ca.) and Digital Research (Monterey, Ca.) and was a Member of the Technical Staff at AT&T Bell Laboratories (Columbus, Ohio). He has degrees from Wheaton College (Illinois), University of California at Santa Cruz, and the University of Texas at Arlington. He is a member of the IEEE and ACM and is a participant in several IETF working groups. He may be reached via e-mail as: `ericf@atc.boeing.com`

[Ed. This article is reprinted from our September 1993 issue].

Announcement and Call for Participation

After the success of the "First International Summer School Advanced Broadband Communications" distributed between Aveiro and Madrid in July 1993, RACE Project BRAIN is pleased to announce its *Second International Summer School on Advanced Broadband Communications* (SS '94) to be held July 11–15, 1994.

Topics

This year the School will be distributed to at least four different and geographically distant sites and will constitute a unique event joining a thorough presentation of ABC (*Advanced Broadband Communications*) with its real use and demonstration. It will include tutorials, in depth lectures, panels and active syndicate sessions covering the most relevant topics of ABC, including:

- Cell based technologies (ATM, Frame Relay, SMDS)
- Access Networks: Mobility, LANs, ATM, ...
- Network interconnection
- Corporate and Virtual Private Networks
- Cost Modeling
- Systems Engineering and ABC
- Management of ABC Systems
- Multimedia Interfaces and Applications
- Entertainment and ABC

Most important, the 1994 Summer School itself will be a demonstration of ATM based broadband communications, applications and services where the results from different RACE projects will be shown in real operation.

Distributed event

The 1994 Summer School will be a distributed event where a multimedia CSCW (*Computer Supported Cooperative Work*) tele-education application will join the lecture rooms of the different physical sites into a unique virtual lecture room such that lecturers and participants loose the sense of physical separation and work together with full interaction. At least, the following sites will join SS '94:

- ETSI Telecomunicacion, Madrid, Spain (Central Site)
- University of Aveiro, Aveiro, Portugal
- CET, Aveiro, Portugal
- TIDSA, Madrid, Spain

Existing trans-national ATM links connecting European Broadband Islands provide the necessary communication infrastructure for SS '94. This will require a complex collaboration between several projects most of them belonging to the RACE program. For example, the ISABEL-IBER projects will provide the RIA (Aveiro) and RECIBA (Madrid) Broadband Islands interconnection, thus supplying the core infrastructure of SS '94. In addition, feasibility studies are underway with RACE Projects CATALYST, EXPLOIT and BETEUS and with the Spanish Tele-education project ETSIT for interconnecting more sites in Switzerland, France, Germany and Spain.

More information

If you are interested in receiving more information please contact:

Dept Ing. Telematica (SS '94)
ETSI Telecomunicacion
E-28040 Madrid
SPAIN
Phone: +34 1 3367332
Fax: +34 1 3367333
E-mail: SS94@dit.upm.es

Ebone continues

At a meeting between the present partners in Ebone on 22 February a proposal worked out by RENATER, ACOnet and ECRC for continuing Ebone after July 1, 1994 was discussed.

Structure

The proposal is for a backbone with nodes in Paris and Vienna, and the initial backbone will have a 512Kbps line from Vienna to Paris and a 1.5Mbps line from Paris to the GIX in Washington. Upgrades are planned as the need arises. Ebone will continue to peer with NSFNET and other US networks and Ebone is negotiating peering agreements with EUnet and EMPB in order to secure good connectivity within Europe.

AUP free

Ebone will remain AUP free, i.e., it will be open to all IP providers. Partners may connect either to Paris or to Vienna and due to the simplified backbone it is expected that the cost will be lower than for the present Ebone.

At the meeting, research networks in nine countries stated their intention to join the new Ebone and others are invited to contact Ebone managements for more information.

Christian Michau, RENATER
Ebone chairman
michau@urec.fr

Frode Greisen, UNI-C
Ebone general manager
frode.greisen@uni-c.dk

[Ed.: See "EBONE, The European Internet Backbone," by Bernhard Stockman in *ConneXions,* Volume 7, No. 5, May 1993, as well as "Global Connectivity: The Global Internet Exchange (GIX)," also by Bernhard Stockman in *ConneXions,* Volume 7, No. 11, November 1993.]

Write to *ConneXions*!

Have a question about your subscription? Are you moving, and need to give us your new address? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Need a *ConneXions* binder for your back issues or perhaps a *ConneXions* sweatshirt? Want to order some back issues? (there are now over 85 to choose from; ask for our free 1987–1992 index booklet and the 1993 index sheet). We want to hear from you. Send your questions, comments or suggestions to:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City, CA 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: connexions@interop.com

CONNEXIONS
303 Vantage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1